


REVISED
EDITION

PRESENTING THE

ATARI® ST™

An in-depth look at the
sensational new computer that
promises to bring you "Power
without the Price"



A Data Becker book published by

Abacus

You Can Count On




Software

PRESENTING the ATARI ST

By L. English and J. Walkowiak

A DATA BECKER BOOK

Published by

Abacus You Can Count On  **Software**

Revised Edition, March 1986
Printed in U.S.A.
Copyright © 1985

Copyright © 1985

Data Becker GmbH
Merowingerstr.30
4000 Dusseldorf, West Germany
Abacus Software, Inc.
P.O. Box 7219
Grand Rapids, MI 49510

This book is copyrighted. No part of this book may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording or otherwise without the prior written permission of Abacus Software or Data Becker, GmbH.

Every effort has been made to insure complete and accurate information concerning the material presented in this book. However Abacus Software can neither guarantee nor be held legally responsible for any mistakes in printing or faulty instructions contained in this book. The authors will always appreciate receiving notice of subsequent mistakes.

ATARI, 520ST, ST, TOS, ST BASIC and ST LOGO are trademarks or registered trademarks of Atari Corp.

GEM, GEM Draw and GEM Write are trademarks or registered trademarks of Digital Research Inc.

IBM is a registered trademark of International Business Machines.

ISBN 0-916439-33-X

Preface

A new era is beginning in the world of computers. The line which separates the home from the business computer is becoming less distinct as machines such as the Atari ST appear. These new super-machines carry a price tag aimed at the home user - especially those who want to upgrade from their present 8-bit home computer. But because of the high performance capabilities of these new computers, they can also be used for serious business, scientific and engineering applications.

Although the Atari ST comes to the market at a price which is extremely attractive, its specifications and capabilities set totally new standards.

- The 16/32 bit Motorola 68000 microprocessor offers very powerful processing capabilities found only in considerably more expensive business and scientific computers.
- The 512K/1M of RAM expandable to 4 Megabytes.
- GEM, the new graphic oriented user interface, enables novices to quickly become productive with the computer.

The ingredients of Atari's ST are already found in other computers. By themselves these features are not so spectacular, the real innovation, however, lies in the low price at which the whole package is offered.

Ford's Model T was no great technical innovation. But the efficient production techniques and the marketing strategy made the Model T very affordable to the consumer. That made the Model T so successful for Ford. This is why Atari describes their ST with *"Power Without the Price."*

Innovation here, represents a leap in the price to performance ratio of home computers. This kind of innovation is a tradition for one of the most significant "movers" in the computer world, Atari Chairman Jack Tramiel.

His first success in home computers was the PET 2001, one of the first out of the box and ready-to-run machines. Its popularity proved so great that enthusiasts had to wait months to buy one, after it was introduced in 1977. But the PET's success was soon overshadowed by that of the VIC-20. The VIC-20 was a smashing success owing to its fine color graphics, sound synthesis capabilities, available peripherals and most of all its affordable price. Millions were sold. To prove that he was not spoiled by success, Tramiel revealed his second masterpiece, the Commodore 64. The '64 remains the most successful home computer to date. More than four million owners are convinced that at such an attractive price, the '64 is *the* computer that meets their needs.

Now Atari is hoping that the ST will be Tramiel's third masterpiece. At the original price of the Commodore 64, the ST offers very exciting performance features of much more expensive computers. To the general public, the ST offers a new spectrum of very ambitious computer applications that were previously unavailable because of price considerations. Home computer users can tap completely new areas with the ST, while business users may quickly find that they can replace considerably more expensive computers with the Atari ST and thereby save a good deal of money too!

The new ST was developed in practically record time by Atari. Our plans for the creation of this book were equally ambitious. Only three months separated the first contact with the ST prototypes and the planned delivery deadline of the book. Despite working night and day, this was possible only because of the in-depth computer knowledge and work experience of our two authors. Lothar English is not only a well-known best-selling author, but has also been a DATA BECKER systems specialist for four years. DATA BECKER is the largest computer software and book publisher in West Germany. He has a detailed knowledge of practically all microcomputer processors and operating systems. Over a year ago he equipped his own computer, a CBM 8296, with an additional board sporting a Motorola 68000, allowing him to prepare for the new world of 16/32 bit processors. Joerg Walkowiak entered the world of data processing in 1978 with his Radio Shack TRS-80. Joerg also is an experienced book author, having written several books on computer subjects.

The goal of this book on the Atari ST is not only to provide the reader with a summary of the capabilities and features of this fascinating new machine. It should also serve as a good source information for prospective buyers. Additionally, we hope that it will suggest potential uses for this powerful new computer.

The information in the first edition of this book was derived from hands-on experience which the authors received during their work with prototypes of the new Atari ST. As you are aware, computer development is a creative and on-going process. For this reason, the information presented in the first edition deviated from the actual production models and the accompanying software. This new revised edition has been updated using actual ST production models and released software, hopefully correcting any misinformation presented in the previous edition.

This edition is divided into two sections. The first section covers what a computer operating system does and why the ST is truly "user-friendly." Then we will briefly discuss the software included with the ST system. The second section describes in greater detail the actual components and operating system of this fantastic new computer.

We wish to thank the Atari Corporation for making the ST available to the authors and for their technical assistance. Good luck, Atari!

TABLE OF CONTENTS

1.	The ST, a user-friendly computer	1
1.1	The Traditional Office	5
1.2	Prerequisites for a friendly computer	7
1.3	The Mouse	9
1.3.1	Working with the Mouse	10
2.	Working with GEM	12
2.1	Menus	17
2.2	Windows under GEM	25
2.3	The Disk Directory	30
2.4	Working with Files under GEM	34
3.	Software	43
3.1	1ST WORD	46
3.2	NEOchrome	49
3.3	DB One	50
4.	Computer Languages	53
4.1	Problem-oriented Programming Languages	58
4.2	Interpretive Languages	60
4.3	ST LOGO	63
4.3.1	Procedures in LOGO	66
4.3.1.1.	LOGO Variables	68
4.3.1.2.	Arithmetic operators in LOGO	69
4.3.1.3.	Logical operators with LOGO	69
4.3.1.4.	Controlling the program execution	70
4.3.1.5.	Graphics commands	71
4.3.1.6.	Primitives for word and list processing	74
4.3.1.7	LOGO Summary	78
4.4.	ST BASIC	79
5.	Introducing the 16-bit processors	83
5.1.	The true 16-bit processors	88
5.2	The 68000 processor	90
5.3	The Instruction set of the 68000	95
5.3.1	8-bit verses 16-bit	101
5.4	Advanced features of the 68000	104
5.4.1	Relocatable programs	106

6.	The Architecture of the ST	107
6.1	ST Interfaces	112
6.1.2	Floppy disk interface	114
6.1.3	The Hard Disk	116
6.1.4	High speed through DMA	118
6.1.5.	The Printer Interface	120
6.1.6.	The Serial Interface	122
6.1.7	The MFP 68901	126
6.2	The sound capabilities of the ST	132
6.2.1	The MIDI Interface	134
6.3	The ST Keyboard	135
6.3.1	The ST Keyboard processor	137
6.4	The ST Video Interface	138
7	The Operating System (TOS)	145
7.1	The BIOS and BDOS	151
7.2	BDOS Functions	153
7.3	The BIOS	159
7.3.1	The XBIOS	161
8.	Inside GEM	165
8.1	Virtual Device Interface	168
8.2	Application Environment Services	170
9.	Conclusion	171

LIST OF FIGURES

1-1: User interface	12
2-1 GEM Window layout	28
2-2 GEM Menu diagram	29
2-3 Disk Information	33
2-4 Install Printer Accessory	40
2-5 Set RS232 Parameters Accessory	41
2-6 Control Box Accessory	42
3-1 DB One Sample Screen	51
4-1 ST BASIC Screen	80
5-1: 68000 Block diagram	88
5-2: 68000 Pin layout	89
5-3: 6502 Register set	91
5-4: 68000 Register set	92
5-5: 68000 Operands	94
6-1: Memory map	110
6-2: I/O Assignments	111
6-3: Block diagram of ST	113
6-4: 3 1/2" Floppy disk	115
6-5: DMA Port	119
6-6: Keyboard layout	136
6-7: Bit mapped graphics	139
6-8: ST Screen photo	142
6-9: ST Screen photo	143

LIST OF TABLES

51-1: The 68000 Instruction set	96
6-1: MFP 68901 Internal registers	127
6-2: MFP 68901 Interrupt priority	129
6-3: Interrupt structure of the ST	131

Chapter 1

The ST, a user-friendly computer

- 1.1 The Traditional Office
- 1.2 Prerequisites for a friendly computer
- 1.3 The Mouse
 - 1.3.1 Working with the Mouse

1. The ST, a user-friendly computer

The ST represents a major technological achievement in the price to performance ratio of computers. The next few pages will explain this achievement in operating system design and what it means to the user. To fully appreciate the significance of the new ST's operating system a little history is required.

Over the years, a large number of operating systems have been developed for computers. In general, the newer systems have included many new features and improvements over their predecessors.

The major tasks of a computer operating system are:

- handling the commands entered by the user;
- loading user programs into the appropriate area of memory in response to the commands;
- giving these user programs access to the computer (execution);
- coordinating several user programs in a multi-user environment;
- managing the memory that is allocated for working storage;
- controlling and monitoring the flow of data among the different peripheral devices;
- resetting the normal state of the computer after the user program ends.


Most of these tasks are performed without the user even noticing that they are taking place. However, one area which has been particularly confusing to many users is the management of the peripheral devices (printers, disk drives, ect.).

Here the user is confronted by the rigid rules and procedures of the operating system. The procedures do work, but often presume that the user has in depth knowledge of the operating system. The user must be familiar with terms such as:

- formatting
- file
- directory
- erase file
- backup or copy
- stat or pip

Terms such as these often frighten or confuse the novice computer user. To make the computer easier and more friendly for the user, operating system designers have tried to simplify the way in which the user enters commands into the computer. The goals of these approaches are to make it possible for the user to perform productive work without extensive training or reference to volumes of technical handbooks. Can you understand this CP/M+ command: `PIP B:=A:* .DOC [G10V]`? This command copies all files with the extension `.DOC` on disk drive A, user area 10, to drive B and then verifies that the copy was successful. The older operating systems were said to be rather cryptic, to say the least.

Drawing from years of in depth study at Xerox Corporation's Palo Alto Research Center, researchers observed that office employees do not want to give up their typewriters, filing cabinets and erasers in exchange for new technology. Instead they want to continue working in a familiar environment. So one approach to computer automation is to simulate this environment and its tools with the computer. The ST's operating system, TOS, does this.



1.1 The Traditional Office

If you think about what tasks a teacher in his study, the secretary in the reception room, the journalist in his open office and the business executive perform, you discover that they all work with similar tools. If we ignore the telephone, the list of tools looks like:

- desk
- typewriter
- blank paper
- correspondence
- file folders and portfolios
- several pens for notes and drawings
- an eraser for corrections
- a trash can for hopeless cases

and more:


- a calculator for basic arithmetic operations
- a clock so that he/she knows when lunch break starts or as a reminder for appointments
- a copy machine so that important papers are not accidentally lost and forgotten.

Let's imagine ourselves at the desk of an office manager. A typical work day begins with the employee sitting down at his clean desk. The first assignment is to produce a sales report for the past month. The manager turns to his typewriter, inserts a piece of paper and types a few introductory lines. Next the manager looks through the sales folder containing last months records and spreads the daily sales logs out over the desk to get a better look at them. Next the manager types the daily sales totals. Then he/she reaches for his adding machine to calculate the monthly total. Everything seems to be going all right until it is discovered that several days sales were overlooked because one sheet of paper obscured another.

So the manager inserts another sheet of paper into the typewriter, retypes everything, but this time not forgetting to include the missing day's sales. The first sheet of paper is thrown into the trash can and the new report sent off to the boss. One further thing: a photocopy of the report is made for the accounting department.

So to make the transition from a traditional office to an automated office as easy as possible, the work process should look the same when the office worker uses the computer. In addition, to increase productivity the worker's interaction with the computer should be easy to learn and understand.

The operating system of the ST makes the transition from the traditional office to a computerized office extremely easy and fast. This is performed by a "user-friendly" operating system. When the ST is started the user is not confronted by a blank screen displaying a cryptic >A:. Instead there is a graphic representation of a **desktop**, complete with accessories, familiar files and commands to open the files.



1.2 Prerequisites for a friendly computer

A "friendly computer" must be blessed with a certain set of technical capabilities. These should allow a user to operate the computer without extensive training in software or electronics.

Due to the rapid advances in the micro-electronics field, new integrated circuits have been developed that have enabled manufacturers to produce very capable computers, which the everyday user can afford. Likewise, software technology has made great strides: software for microcomputers has become more reliable, affordable and creative.


This advanced technology has enabled ATARI to create the ST with a resolution of 640x400 points at an affordable price. This high resolution is suitable for representing familiar objects on the screen -- objects such as a file cabinet, sheets of paper or even a trash can. These objects have come to be known as *icons*.

Icons allow a first time computer user to be immediately familiar with the computer. Think about how you would copy a folder full of papers. You would first get the folder containing the originals from a file cabinet, get a new folder, go to a copy machine, make the copies and then put these copies into the new folder.

In the picture on the next page notice the trash can icon. It looks like a trash can! The floppy disk, where you store your documents, is represented by a file cabinet drawer. The open file drawer is represented by the box in the middle of the screen. It even contains the familiar manila envelopes!

The icons used by the ST are immediately familiar to the first time user. The folder appears on the screen, a new folder may be created, and the required documents can then be copied into the new folder.

The folder appears on the computer screen using a software innovation known as *windowing*. On a computer screen, windows simulate several file drawers or sheets of "paper" -- one on top of another. The user can "flip" through these drawers and display a specific page by selecting the appropriate window.



To be able to easily select a specific work task, or to select from among the different windows, a hardware device called a *mouse* is used. A mouse can be thought of as an extension of your hand. You can not reach into a computer to select a folder as you would a file cabinet. The mouse allows you to select a folder or document on the computer screen.


It is the use of this mouse as an input device that the ST hopes to stir up the home computer market. While it is true that computers such as the Apple Macintosh use these same techniques, Atari has been able to do the same at a very affordable price. No longer is this exciting computer technology the privilege of a few. Instead the ST makes it possible for the masses to afford a computer that can be operated easily and efficiently through graphic oriented displays.

1.3 The Mouse

The mouse is a small input device, that houses a small ball whose rotation in two axes can be measured, either mechanically or optically. If you turn a mouse on its back, you can see that it's a variation of the more familiar trackball used in many video games. With a trackball, the user rotates the ball in the desired direction with short movements of the palm of the hand. With a mouse, the user moves it in a desired direction over a flat surface.

The operating principle of a mouse is based on a coordinate system where a movement in a given direction represents a given **x** and **y** value. If the flat surface over which the mouse is moved is said to correspond to the computer's display screen, then each point on the screen can be assigned an x-y coordinate pair. Thus a movement by the mouse represents a corresponding movement on the screen.

If software is written to support a mouse, then it's possible to replace the keyboard for many tasks. One method is to display a menu of choices on the screen. The mouse is used to position an arrow or cursor to the desired choice and a button on the mouse is pressed. This is called *clicking* the mouse. So you can select a function with two simple steps: moving the mouse and clicking it. You might even think of the mouse as an extension of your finger, where you point to the symbols displayed on the screen. The Atari ST is designed to be used with a mouse.



1.3.1 Working with the Mouse


How is the mouse used to perform various tasks?

The ST's impressive hardware is not very useful without appropriate software. Let's see how we can do productive work with the ST.

Before using a blank diskette, it must be formatted. One of the functions of the TOS operating system is formatting a diskette.

With most other computers, the user must know the name of the command (FORMAT) to perform a desired function. As is the case with many of these commands, one or more parameters must be entered as part of the command. So the user must also know how many parameters are required and the order and syntax of these parameters. Only then can he enter the full command at the keyboard.

One alternative to entering the command at the keyboard is to display a list of functions on the screen. Then using the mouse you can click the desired function. The user needs only remember the correct command, but not a complicated syntax. Although simpler to perform than typing the command, the method is a long way from the automated office.



Now let's introduce a slightly different concept. In place of the list of functions, why not represent them as icons (graphic symbols) on the screen? The ST does just this by using an operating system called GEM. GEM (Graphics Environment Manager) is a graphics oriented operating system from Digital Research. In practice, GEM is not really a full operating system. Rather it is an interface between the user and the operating system. It's purpose is to simplify the communication between the user and the actual operating system called TOS. Figure 1-1 illustrates how the hardware and software components of the ST fit together.

At the lowest level is the computer hardware. It is not directly controllable by the user. Instead, the operating system at the next level controls the hardware. The user at level 3 communicates with the operating system through programming languages, utilities or application programs at level 2.

By introducing another level of communication we can reduce the amount of knowledge that the user requires in order to interact with the language, utilities or application programs. So move the user interface up to level 4 and place GEM at level 3. Now, we have a very easy way to use the ST through the GEM interface. No longer does the user have to remember large amounts of cryptic commands or continually page through massive technical manuals to perform simple functions.

This user interface greatly reduces the training necessary to use the computer. It also reduces the apprehension felt by first time computer users since the icons represent objects they are familiar with.






Figure. 1-1 User Interface




Chapter 2

Working with GEM

- 2.1 Menus
- 2.2 Windows under GEM
- 2.3 The Disk Directory
- 2.4 Working with Files under GEM

2. Working with GEM

After turning on the ST, the screen displays a pictorial representation of a desktop. It is divided into the working surface of the user and the system menu header.




Icons appear on the desktop along the left side of the screen. These icons are figures of the peripheral devices connected to the ST. For example, two light colored file drawers represent the two ST floppy disk drives - one called A:\ and the other B:\. Below the disk drive icons is an icon of a trash can.

A small black arrow appears in the middle of the screen representing the position of the mouse. This is called the mouse pointer or arrow. The mouse is the main input device that the ST uses under GEM.

There are two fundamental operations that can be performed with a mouse: selection and dragging.


To select an object that appears on the screen, simply move the arrow across the screen with the mouse until it is positioned over the desired icon and then click the mouse button. The selected icon is then displayed in reverse on the screen, in black. This indicates that it is selected.

To move an object from one location to another on the screen, you must drag it. To drag an object, move the mouse pointer across the screen with the mouse until it is positioned over the desired icon, press and continue to hold the mouse button, and then reposition the icon on the screen by moving the mouse. When the icon is in the desired position, release the mouse button. As long as you continue to press the button, the icon can be moved from place to place on the screen.




2.1 Menus

As already mentioned, GEM uses both icons and menus. The use of menus allows a user to quickly become familiar with the operation of the computer without having to read volumes of training manuals. When the ST is first turned on the following menu is displayed at the top of the screen:




Now if you position the arrow over one of the menu choices a *pull-down menu* will appear. A pull-down menu works like a common household window-shade. The shade may be pulled down when needed or retracted if not needed. If needed, a pull-down menu displays a further set of choices.


So by positioning the arrow over the **Desk** menu the following pull-down menu appears and displays a set of functions that you may use on your automated desktop:



The pull down menu **File** contains all the functions for manipulating files, whether the file is a program, document, or even the data disk:



The pull-down menu for **View** displays the contents of ST's disks, represented by file drawers. You can display the files in a variety of ways. The contents may be displayed as icons or as text. In addition, they may be arranged by name, date, size or type. The pull-down menu looks like this:




The pull-down menu for **Options** allows you to set many system-specific parameters without consulting technical manuals:



The install functions are used to adapt peripheral devices or other programs to GEM. You may install different disk drives, such as a hard disk or even a super fast RAM disk. A RAM disk is resident in memory and emulates a floppy disk. A RAM disk is super fast because memory to memory operations are much faster than memory to peripheral operations.

Install Application allows you to let applications (programs) automatically start from their data files. Let's say you use a word processor program, you can simply select a document that was created by the word processor and it will start the word processor with the selected document ready for editing or printing. Older operating systems required you to first load the word processor and then load the document to be edited.

System parameters can be changed with **Set Preferences**, this is for the experienced ST user. GEM is designed to be as user-friendly and forgiving as possible. Therefore, confirmation is requested before a copy or deletion is performed. For example, when deleting a file, a *dialog box* appears on the screen requesting confirmation.




At this point, the user can cancel the operation by clicking **CANCEL** on the dialog box. The **Set Preferences** choice can be used to alter this confirmation process. It is also used to set the screen resolution to Low, Medium or High.


By using these four pull-down menus and the icons, you can perform all of the usual computer operations. These include formatting diskettes, deleting or copying files, viewing the directory of a disk and starting applications. All of these operations can be performed without ever once having to use the keyboard or remembering cryptic command sequences.

Here's an example: Select a disk drive by positioning the arrow over the file drawer icon and press the mouse button. The file drawer icon now appears in reverse indicating that it has been selected.

Suppose you want to know which files are on the disk in this drive. Position the arrow to **File** and a pull-down menu appears. Drag the arrow to the **Open** choice and press the mouse button. Immediately, the pull-down menu disappears and a new window appears (the window is said to have opened) displaying the contents of the diskette on the screen.



To show the contents as text instead of icons, position the arrow to **View** and another pull-down menu appears. Drag the arrow to **Show as Text** and press the mouse button. Again, the pull-down menu disappears and the contents of the window is rearranged. The disk files are now displayed as text instead of icons.




Now let's take a closer look at windows. A GEM window is made up of several elements: a border surrounds a light background; arrows are in the corners of the windows; a black rectangle is found in the upper left and lower right corners.

In the previous figure, you can see that a section of the border is cross-hatched. This means that only part of the available information can be displayed within the window. To display the rest of the information, you can scroll the window using the scroll arrows.

The arrows are used to *scroll* the contents of the window. If when displaying information on the screen, it will not all fit within the window dimensions, you can use the arrows to scroll the information within the window. To scroll downwards, for example, you would position the mouse arrow over the down scroll arrow located along the right-hand edge of the window and click the mouse button. The information on the screen will then scroll downwards.

To display more information on the screen, you can also change the size of the window. Do this by dragging the window's size-box located in the

lower right-hand corner of the window to a new location. As you drag the size-box, you can see how the dimensions of the window change. The maximum size of a window is determined only by the size of the work surface available. The window may be enlarged to cover even the TRASH icon for example.




As you can see, the menu titles remain displayed at the top of the screen.

If you have completed using a window, you should **Close** it. Do this by clicking the close box in the upper left-hand corner of the window or selecting **Close** from the **File** menu.. The window disappears and any information that was hidden by the window is again visible.


While all of this is taking place on the screen, the processor is performing many tasks. In order to restore a screen when a window is closed requires that the previous contents of the screen be saved. If multiple windows are superimposed on the screen, they too must be saved. To do this the operating system must have a lot of memory available to use and must work quickly since the user does not want to wait long for the new screens. The 68000 processor in the ST is ideally suited to this task.

In the picture below notice that the icon for a file called **READ.ME** appears in reverse, meaning that it is selected. Instead of selecting an icon and then going to the **File** menu to **Open** a file, you can position the mouse arrow to an icon and click the mouse twice in quick succession (called *double clicking*). Double clicking is a shortcut to opening a file. By double clicking the icon for **READ.ME**, we can open the file.



2.2 Windows under GEM


A window consists of a border surrounding the viewing field. At the top of the window is the title box displaying the name of the file or disk. Beneath the title box is a rectangle displaying the amount of storage space that the files contain. It looks like this:



A window currently in use contains a close box in the upper left hand corner. A window that is overlapped by another does not have a close box.

On occasion you might want to enlarge a window to make as many objects as possible visible at one time. On the other hand, you might want to reduce the window size in order to see a partially covered window underneath. To change the size of the window opening, drag the size box, located in the lower right hand corner, until the window is of the desired dimensions and then release the mouse button.

You can also relocate an entire window on the screen. To do this, drag the title line of the window to the new location and then release the mouse button. You'll note that as the window is being dragged to a new location, only an outline of the window is moved. The window itself is not relocated until the mouse button is released.



At any given time, four windows may be opened but only one window may be *active*. To make a window active, you simply move the mouse arrow to the window and press the mouse button. GEM moves this window on top of any other windows currently open. When you complete your work on this window, you can activate another window in the same way.

Desk File View Options




Figure 2-1 GEM Window Layout





Figure 2-2 GEM Menu Diagram

Desk	File	View	Options
Desktop Info... VT52 Emulator Control Panel Set RS232 Config. Install Printer	Open Show New Folder... Close Close Window... Format...	Show As Icons Show as Text Sort by Name Sort by Date Sort by Size Sort by Type	Install Disk Drive... Install Application Set Preferences... Save Desktop Print Screen

2.3 The Disk Directory

The disk directory can be displayed in several different formats by using the **View** option from the desktop menus. By default, the directory is displayed as icons as in the following picture.



The amount of disk used is displayed in addition to the file names and types.

Among the different file types are:


folders



programs




data




The icon for a data file is a sheet of paper. The icon for a program is a small window.

The icon for a folder is a symbol that looks like a familiar manila file folder. The folder icon is a way to group an arbitrary number of program or data files together. By grouping these files in a folder, the user can easily organize the information on the disk. To see the contents of a folder, you merely open the folder (by either double clicking or by selecting and then OPENing it.)

As an alternative to displaying the directory as icons, you may also display the contents alphabetically by file name; alphabetically by file type; chronologically by creation date; or by file size.




You can select individual files by using the mouse pointer. In addition to information about the individual files or the folders, you can also call up information on the status of a diskette or a file. The operation **Show Info** in the **File** menu is used for this purpose.



In the next example, drive A: contains no data at all. You should notice that the number of files and folders, the amount of unused space on the disk and the name of the diskette are displayed.

After viewing this information, click the **OK** field in the dialog box to return to the previous window. Notice that the **OK** field is outlined with a heavy line. This indicates that you may also simply press the <RETURN> key to return to the previous window.


Figure 2-3 Disk Information



2.4 Working with Files under GEM


Backing up files is one of the most fundamental file operations. Using GEM, it is also one of the easiest operations to perform. Let's see how we would back up a file using the ST. The source files are on drive A: and we will copy them onto a disk in drive B:.

The first step is to select drive B:. Do this by clicking the icon labeled **Floppy Disk B:**. If the diskette in this drive is blank, format it by selecting the **File** menu and dragging the arrow until the **Format** operation is selected then press the mouse button. Since GEM is user-friendly you'll see the following message on the screen:



If you've made a mistake, you can leave this procedure by moving the mouse pointer to CANCEL. By clicking OK the formatting begins.

When the light on the drive goes out, we're ready for the next step.




To copy a file, simply drag the icon for the the desired file until it is positioned over the icon of the destination drive. When the destination drive is darkened release the mouse button. Copying begins.


You can also manipulate multiple files as a group. To do this, the icons for the individual files are selected as a group. To select the group, it's necessary to form an imaginary rectangle around the desired icons. First, position the mouse arrow to one of the two upper corners of the imaginary rectangle and press the mouse button. Now drag the arrow to the opposite corner of the rectangle. As you are dragging the arrow, a dashed line appears on the screen to show you the icons which have been selected. When all of the desired icons have been selected (are within the dashed rectangle) release the mouse button. All selected icons will now appear in reverse.

The selection of a group of icons can also be performed by selecting the individual icons with the mouse while holding down the shift key.

By moving one of the group of selected icons, all of the icons appear to move together. The selected group of files are displayed as outlined icons as they are dragged. To copy these files to a backup disk, drag the icon of one of the group of files until it is positioned over the destination disk drive and release the mouse button. The selected files are copied to the backup disk one at a time.




This method of manipulating files is very easy to learn and very powerful as well. It can prove dangerous too since it's possible to toss many files at a time into the **Trash** can. You'll recall that GEM is very user friendly and before files are deleted, accidentally or intentionally, the following dialog box appears.





So we've seen a very convenient way to group files together. Now let's look at another way. Go to the **File** option, drag the mouse arrow to **New Folder** and press the mouse button. A dialog box appears on the screen. Now you'll have to use the keyboard to enter a name for the file folder. A folder name consists of up to eight characters with an optional three character extension. The extension is used to indicate the type of file, .PRG is a program file, .DOC is usually a document file. The extension .TTP is a special type of program file. The TTP stands for TOS Takes Parameters and allows you to enter data that the program may require. When you have typed in a file folder name, click the **OK** box.

To place a file into this folder, drag its icon until it is positioned over the folder and then release the mouse button. The file is copied into this folder. Naturally, you can copy as many documents to the file folder or create as many folders as necessary to achieve your desired degree of organization. You now have two copies of the same file on the disk, one in the folder and the other in the desktop. Two copies of the same file may seem redundant but it insures the safety of your files. This method also gives TOS compatibility to MS/DOS (IBM-PC) file and directory (folders in TOS) structure. To save disk space you may delete the original file on the desktop by moving it into the **Trash Can**.



The final example, is the copying of entire diskettes. This is how easy it is to copy an entire diskette. Use the mouse pointer to select the source disk drive, drag the icon until it positioned over the destination drive and release the mouse button. You'll see the following messages appear on the screen:





Before leaving our discussion of GEM, we want to quickly point out a few other features.

The Desk menu contains accessories that can be accessed at most times. These include a VT52 terminal emulator, a telecommunication program that allows you to talk to other computers with a modem. The Control Panel allows you to set the time and date, alter the keyboard and mouse response, set the colors and turn off audio feedback. You may also set the RS232 and printer parameters quite easily from this menu.

These simple tasks were quite complicated on older computer operating systems and could not be accessed all of the time. All of these tasks were usually performed manually when the computer was first turned on.

The term accessories is used because the ST is not limited to only the present choices in the **Desk** menu. Other accessories can be added to the ST. Useful desk accessories that will be added are a calculator, easier printer controls and possibly simple games.

Figure 2-4 Install Printer Accessory




Figure 2-5 Set RS232 Parameters Accessory





Figure 2-6 Control Box Accessory



Chapter 3

Software

- 3.1 1ST_WORD**
- 3.2 NEÖchrome**
- 3.3 DB One**

3. Software

Software is used to describe the programs that allow the computer to do useful tasks. Without software a computer can not do anything. Languages are programs that allow users to write software. The ST series of computers will have a wide variety of both thereby assuring its popularity.

With the introduction of the ST, Atari announced the release of three very powerful programs, a powerful word processor, database and a color art program. These programs are available from the Atari dealer where you purchased your ST. The ST also has two very powerful programming languages, ST-BASIC and Atari LOGO. The advanced capabilities of the ST assure a wide variety of software releases, check with your local Atari dealer for the current list of available titles.

In this chapter we will give a brief overview of the three programs mentioned above and in the next chapter the languages since they are both considered software. A computer is only as good as its software and these first software releases show only a fraction of the capabilities of the ST.

Programming languages are important because they insure that software will be written to use all of the features of the ST. The ST currently has some of the best programming languages available. The wide variety of programming languages will insure a vast selection of quality software.

3.1 1ST_WORD

Atari first introduced ST Writer which is a powerful word processing program based on the best selling AtariWriter program that ran on the Atari 8-bit computer systems. This is a very capable word processing program but it does not use the GEM interface available on the ST. It is awkward to use for first time users because of the large amount of commands they had to remember.

To take full advantage of the GEM interface Atari released 1ST_WORD which makes use of GEM. This program is perfect for the first time and experience word processing users. The menus and commands are always available to the user. You do not have to memorize a complex series of commands to use the program. They can all be called directly from the screen. For first time users the tutorial introduces you to the basic elements of word processing. These include saving and loading files, searching and replacing text, deleting individual letters and lines of text, formatting text and various printing functions.

The documentation is included on the disk, all you have to do is print it out on your printer to have a complete manual. The program is set up for standard and Epson compatible printers. You can also configure 1ST_WORD to work with a variety of printers. Ask your dealer to configure this program to your printer if you are unsure of the method explained in the manual.


One of 1ST_WORD's most powerful features is its ease of use. The text appears on the screen as it will appear when printed. With many older word processors you never knew exactly what the text would look like until it was printed. Different print styles, such as **bold** and underlined are available to spruce up your letters or documents.

1ST_WORD is a very capable word processor at a price that can't be beat. However it does not use all of the tremendous capabilities of this computer. Future word processing releases from a variety of other software publishers will undoubtedly incorporate many new features previously available only on very expensive typesetting machines.


The figure below is an actual screen dump of 1ST_WORD. The text on the screen appears exactly as it will be printed. We have included the various font styles, justification types, and different margin settings. All of these functions are selected from the menus displayed on the top of the screen. No longer will the user have to remember sequences of control characters or guess at what the final printout will really look like. The different font, justification and margin settings are simply selected from one of the menus. What could be easier.

1ST WORD also includes a unique help function for first time users. When this function is selected, from a menu of course, a dialog box appears when each menu item is selected. This dialog box explains the function of the selection. First time users will be word processing in record time, without having to consult large technical manuals or remember cryptic command sequences.

On the next page is the actual print out of the document in the 1ST_WORD window and the PRINT FILE screen.



This is bold
Underlined and bold or italic
This is a combination of many styles
 Superscript or subscript can also be
 used.




3.2 NEOchrome

This color graphics program is an preview release version of a complete color graphic designing program to be offered from Atari. NEOchrome has a resolution of 320x200 with 16 colors, so it can only be used with a color monitor. The preview version is quite impressive although it does not completely use the GEM interface. The final release will be a super program if this pre-release is any indication.

If you have any doubts about the mouse as an input device, you must try this program. A short session will soon convince you of the ease and functionality of the mouse .

Like 1ST WORD the documentation is included on the disk. Below is a sample printout on an Epson printer.



3.3 DB One

DB One, from STONEWARE®, is a simple yet effective database. A database can be thought of as a file cabinet in which you place records. The difference between a database and a file cabinet is in the speed of access to the data in each. In a file cabinet you would have to manually resort the records to put them in a new order, say from alphabetical order to zip code order in a mailing list. This could be quite a task with a file cabinet, but with a computer and DB One this becomes very easy. A database also allows you to sort your file in any number of ways, a file cabinet would require manual sorting for each method.

DB One uses the GEM interface so it is extremely easy to operate. It also includes templates for mailing lists, checkbook and date fields. To design a database you simply insert the sizes of your fields on to the screen, a field could be named **LAST NAME**, consisting of 30 characters. These fields can be placed anywhere on the screen to suit your own individual needs.

DB One's templates are very easy to use. You select the spot on the screen where you wish them to appear and then *SPLAT* them into that position from a menu. The term *SPLAT* is what DB One calls the operation of placing the templates into your data base design.

Once you have designed your database the data can be sorted in a number of different ways, depending on the situation required. The data can also be output to a printer so the data in your database can be used for reports or printing mailing labels.

Sorting, creating, searching, updating the database are operations that are performed from the GEM menus. With a little introduction even a first time user will become a proficient database user in an extremely short period.

On the next page is a screen dump of DB One with a sample mailing list being constructed.

Figure 3-1 DB One Sample Screen

Desk

File

Edit

Splat

Options

Organization

First Name1

Address35

City22

ATARI ST Owner3

Comments560

Phone10

Name20

State5

Zip Code9

Date of Purchase15

ign Form

Phone10

Name20

State5

Zip Code9

Date of Purchase15

Mailing List

Checkbook

Collection

Date

Wart Hogs

ZAP

HELP

Chapter 4

Computer Languages

- 4.1 Problem-oriented Programming Languages
- 4.2 Interpretive Languages
- 4.3 ST LOGO
 - 4.3.1 Procedures in LOGO
 - 4.3.1.1 LOGO Variables
 - 4.3.1.2 Arithmetic operators in LOGO
 - 4.3.1.3 Logical operators with LOGO
 - 4.3.1.4 Controlling the program execution
 - 4.3.1.5 Graphics commands
 - 4.3.1.6 Primitives for word and list processing
 - 4.3.1.7 LOGO Summary
- 4.4 ST BASIC

4. Computer Languages

There are a wide variety of computer languages. Hopefully, each different language serves a certain user need. In this chapter we'll describe a few requirements and examples of what different computer languages can do. Then we'll give you a brief overview of the two excellent languages included with your ST system.

A computer language must first of all be unambiguous. Each individual language element must have an exact defined syntax and meaning; ambiguous interpretations are not allowed.

To the casual user, the computer is an intelligent machine. But the fact is, even the fastest, most expensive computer is only a simple machine. Without the correct instructions, it is incapable of performing any useful work.

This short example shows that even a simple addition requires a program to tell the computer how to proceed:

1. Get the first number
2. Get a second number and add it to the first
3. Display the sum of these two numbers

Even though it's easier to perform the addition in your head or on a pocket calculator, here's a sample solution on a computer. We must know some facts before we proceed, for example, from where do we get the two numbers?

Two memory locations may be designated as the source of these numbers. The instruction sequence might look like this:

1. Load the first number into the accumulator from memory location 1.
2. Get the second number from memory location 2 and add it to the number in the accumulator.
3. Write the sum (the number in the accumulator) to memory location 3.

The 68000 processor could follow this algorithm, but not before we have made sure that the two numbers are actually in the appropriate memory locations. Furthermore, we must be able to examine the contents of location 3 to find the result of the calculation.

You might want to write this program in machine language or assembly language. Here you would use the instruction set of the processor which we introduced earlier.

If you take a look at an assembled program, you'll see a series of numbers in the printed listing. These numbers are the hexadecimal representation of the machine language. If you can imagine the hexadecimal numbers as a bit pattern containing only 0's and 1's, then you'll see the actual machine language program as the 68000 processor sees it.

ASSEMBLED 68000 PROGRAM

HEX CODE	SOURCE CODE
1 00000000	.text
2	* This program adds the first
3	* three integers and stores the
4	* result in memory
5 00000000 303900000000	start: move.w a,d0
6	* Load first number
7 00000006 D07900000002	add.w b,d0
8 0000000C D07900000004	add.w c,d0
9 00000012 33C00000000A	move.w d0,d
10	* Store answer
11 00000018 4E75	rts * Return
12 00000000	.data
13 00000000 0001	a: .dc.w 1 * Numbers to add
14 00000002 0002	b: .dc.w 2
15 00000004 0003	c: .dc.w 3
16 00000006 0000	d: .dc.w 0 * Answer here
17 00000008	.end

It is a series of 0's and 1's that the processor receives during execution. The number system comprising of only a 0 and a 1 is known as the binary number system. Suppose for a moment, that you had to communicate with the 68000 by using only these binary numbers. You would have to know the series of numbers for each different instruction and also the series of numbers for the location of the data which the processor required. Needless to say, programming a processor using binary machine code would be very error prone and time consuming.

Assembly language is a great step forward in programming. The bit patterns (binary numbers) for each instruction are replaced by easy to remember *mnemonic* abbreviations. It's easier to remember a mnemonic such as LD (for Load) or LDA (for Load Accumulator) instead of a cryptic bit pattern.

To summarize machine and assembly language programming:

1. The instructions perform relatively fundamental operations such as data manipulation and basic arithmetic. For this reason a program must be broken down into a large number of individual instructions. In doing so, the program become very long. Increasing the number of instructions also increases the probability of errors.
2. Machine or assembly language programs are rather difficult to read and change. Therefore program maintenance is slow and expensive.
3. The mnemonic instructions depends on the processor used and architecture of the computer system. Exchanging programs between different computers is very costly or impractical.
4. The programmer must be very familiar with the operation of the computer and all peripheral devices. He must know how to access, the keyboard, screen, printer, disk drive, etc.

We don't want to downplay the use of assembly language programming. An assembly language program can take full advantage of the processor. Each of the instructions is available to the assembly language programmer where he may hand-tailor the software for highest execution speed. He has total control over the memory used for the program and data.

4.1 Problem-oriented Programming Languages

As computers became more widely used, the base of people using them became wider as well. Not only were the physicist and mathematicians using the computer but others from other disciplines were also using them.

The physicist and mathematician were fairly well-suited to learning computer programming. They found the transition from hand-written algorithms to computer algorithms rather straight-forward.

Others, however, found the transition much more difficult. Furthermore, it was too expensive to describe a problem to a programmer so he could write a program to solve it. For this reason, new artificial languages were developed so that the user could program a solution to his problem in a more familiar terminology.

The first major language to appear was called **FORTRAN**, for **FOR**Mula **TRAN**slator. It was highly mathematical in nature and was developed mainly to suit the mathematical needs of scientists and engineers.

Many other languages followed, a few of which came to fame and wide-spread use, while others which ran on only a few individual computers lasted only a short time.

Languages such as **COBOL**, the **CO**mmon **B**usiness **O**riented **L**anguage, were written especially for use in business data processing.

Probably the most popular language for the ST will be the programming language **C**. **GEM** is written in **C**. The reason that this language will become popular is that the source code is easy to read and is a very portable. By portable we mean that a program written on one computer can be converted to run on another computer very easily. Atari recognized the importance of this feature when designing the ST, more about this in the hardware sections of this book. This feature will make vast amounts of software available quickly. The Atari ST programmers development kit from Atari includes the **C** programming language and an assembler.

These three languages were designed as *compiled languages*. The compiler checked the language syntax and grammar for correctness and converted the individual statements into equivalent machine language instructions. Finally, this resulting program was executed.

The advantages of this procedure are obvious:

- * The programmer can describe his problem in terms that he is familiar.
- * The compiler can translate these terms into the machine language of the computer. In fact, a compiler can be written to translate the problem oriented language into the machine language of any computer. In this way, the compiler and not the programmer need be concerned with the hardware dependent features of the computer.
- * The program runs or executes quickly since it is now translated to machine language.

There are of course disadvantages to compiled languages:

Compiling a program is a multi-step process. The user-written program (called the source program) is entered and written as a file. This source file is then passed to the compiler.

If the compiler detects an error in the source program - perhaps a syntactical error - then the source program must be corrected and recompiled.

If after a program is compiled, an error is discovered during execution, the source program must be corrected and recompiled.

During the writing and testing of these programs, many hours might elapse going through these procedures. One alternative to the compiled programs was the development of interpreted languages.

4.2 Interpretive Languages

An interpretive language is one where the language elements are translated immediately into the machine language of the computer and executed. If a statement like `PRINT 3+5` is entered, the language interpreter would first check the syntax of the expression and then carry out the command step by step.

The command `PRINT` tells the interpreter that something must be displayed on the screen. If the something is a constant (literal), the interpreter knows that a quotation mark must follow the `PRINT` keyword. But if a mathematical expression follows the keyword, then the quotation mark will not be there.

The above example causes the interpreter to first calculate the sum of 3 and 5, and then display the result on the screen. The interpreter then proceeds to process any subsequent commands.

By using an interpreted language, the programmer can check his work at any time. This is in contrast to compiled languages where he has to first write the source file, compile it and then test it.

Because of the way that an interpreted language works, an interpreted program can never be executed as quickly as a compiled program can. But an interpreted language is considered a better alternative for the novice program developer.

His choice of language is based on a combination of the elements which a given language provides and the ease with which he can communicate with the computer.

We have already seen that a number of programming languages have been developed in the course of the last twenty years. The purpose of a new language is either to open a new area of application for the computer, or to make the computer and its language more human.

There are languages, which are used to easily control machine tool operations. It would be impossible to write a bookkeeping package or even a video game in such a language, however.

The development of an all-purpose language was necessary for ease of learning, so BASIC, developed in 1964 at Dartmouth College, allowed beginners an easy introduction to programming. BASIC stands for **B**eginner's **A**ll-purpose **S**ymbolic **I**nstruction **C**ode, which was an accurate description back when the set of commands consisted of only a few dozen instructions. But today, where commands are necessary for controlling graphics and sound processors, in addition to various methods of accessing files and even instructions for creating structured programs, the name "beginner's language" is no longer appropriate.

The ST comes with two interpreted languages, ST BASIC and ST LOGO. Both of these languages have been released in many different forms. The versions that come with the ST are some of the most advanced and easy to use implementations of these languages.

FORTRAN, BASIC and all other languages developed before 1971 are line-oriented programming languages. Each language statement is written on an individual line with an identifying line number. To use a routine of a program from more than one place in that program, you could use a GOTO command to the desired line number. As many programming textbooks maintain, the result was always something called "spaghetti code" - a very complex and chaotic execution path that makes editing and correction very difficult, if not impossible.

Niklaus Wirth put an end to this by developing the programming language Pascal as a structured, easily readable and manageable language. He argued against "spontaneous" program development. In the Pascal language, the programmer must first define the variables and subroutines he will use. Pascal has been very successful as a teaching language at many American and European universities and schools, but has not met with the same success in industry.

The primary reason that Pascal has become so popular is that it has become available for almost every microcomputer. Likewise, the language Forth has also enjoyed wide success on microcomputers.

Forth and Pascal are already available for the ST. See the following picture that displays a Forth's vocabulary. The standard vocabulary is supplemented with words like `PIECHART` and `BARChart` for easy implementation of business graphics. Other extensions such as `SETFrequency`, `TEMPO`, and `BANJO` make music and sound synthesis extremely simple.

```

TASK FORTH ;CODE MULTITASK PEN HEADING TV TX RMARGIN
SECTOR TRACK ADDRESS PAUSE KILL START FIND-TASK ,TASKS
TASK_LIST MAX_TASKS WRITE READ FORMATHD MODE10 MODE5 MODEX
BMODE10 BMODE5 RSENSE HD0; ENDCMD RELQ SEL NKINSKI BARBIE
PAINTPIC GRAYSCALE PAINTLINE PICPAG PICK BARCHART BARVAL
BLTBAR STARTBAR PIECHART PIEDELAY SHOWIN SHOWOUT SHNXT
MAKETHEM SHIFTCOPY MAKECOPY PUTLO PUTHI PIE COPYBLOCK RSHIFT
TEXTCOLOR WPIR SQUIRAL INCIT HOME SETPOS PENDOWN PENUP LEFT
RIGHT SETHEADING BACKWARD FORWARD BCK FND SETH :: MYSELF
IKBD ASCII DSK QWERTY FONT GX9 ATARI PSGR PSWN PSGA RS
WS RT HT FORMAT MAKETRACK MAKESECTOR DATA TAIL HEAD FILL4E
STUFFN PRELOOP STUFFB CHAP0 CHAP! PSW0 SUPERM USERM SSP0
SSP! USP0 USP! CL5 CURSOR-OFF CURSOR-ON WIPE DSPADR? DSPADR!
SETCOLOR RND RND0 FCIRCLE (FCIRCLE) CIRCLE (CIRCLE) DRAWLINE
CLINE PLOT HIRES SURE 40COL 80COL TEXTCLS SCALE MIDIDLY
PROGRAM PLAYBACK RECORD RECORDM OLDCLOCK PTR TUNE PREVIN
EVITA BANJO SANJOSE ROMPLAY PREVIN2 EVITA2 BANJO2 SANJOSE2
TUNE2 SILENCE (PLAYD2) (PLAYD) MIDIPAGE (PLAY) TEMPO SETFREQ
PRESCALE SIGTUNE SILENCE ASSIGN NOFF NON MIDI KEYUP NOTE
MIDIQ GM PM GETAUX PUTAUX IOB0 IOB! TIME? TIME SETTIME
JIFFIES! JIFFIESD DELAY SYSTM TIMER TICKS! TICKSD IOS SIDE1
SIDE0 SEEK D1: D0: DESELECT DRIVE1 DRIVE0 ?DRESET LASTFLOPPY
?DONE DRESET RDC HDC RDL HDL DMA FLPC DMA? DMA0 DMA!
LOW HIGH DMAH DMAH GPIIP Ok fine

```

These high-level languages can be divided into problem-oriented languages and procedure-oriented languages. Languages belonging to the first category give the programmer the means by which to solve special problems in terms that are familiar to the user. Members of the second category cover a broader spectrum of applications. Where programs are made up of a series of incremental procedures that solve the problem.

In the past few years, another group of languages have emerged which are growing in importance. These are the list-oriented programming languages; languages that work less with numbers and more with objects.

The most well-known of these languages is LISP for LISt Processor. LISP has been used especially in the area of artificial intelligence. This branch of computer science is concerned with describing human models of behavior to computers, to make them think and act like humans.

4.3 ST LOGO

Logo works with lists and is quite similar to LISP, which is why it is described as a subset of LISP by many researchers. As with Forth, it works with a stack and can be expanded with user-defined keywords. Like Pascal, procedures can be defined and variables can be declared as either global or local.

It's greatest advantage is the ease at which you can learn it. You can work with Logo after learning just three commands. You do not need to have a thorough overview of language to use it. Logo is therefore an ideal language for beginners although this does not mean that complex applications are not possible with Logo!

Digital Research has created a version of Logo for the IBM PC. St Logo is very similar to Dr. Logo and is one of the standard programs of the Atari ST computers!

ST Logo represents the most comprehensive implementation of this language available up to this time.

In the late 1970's, Logo caught the attention of the public with its turtle graphics. It was at this time that Dr. Seymour Pappert at the Massachusetts Institute of Technology developed turtle graphics to teach children about computers. Not only was he able to teach five and six-year olds about computers, but turtle graphics also enabled them to program the computer. He used the help of a game to win the interest of these children?


A child places himself in the middle of a room, and the other participants in classroom direct him using verbal commands so that he can trace mathematical figures with his movements. A square might look like this:

```
go forward five steps
turn one quarter turn right
go forward five steps
turn one quarter turn right
go forward five steps
turn one quarter turn right
go forward five steps
```

By using a piece of chalk to mark the path he travels, a square is drawn.

The result of the experiment can be discussed with the children and additional "games" of this type can follow. So the children can learn geometry by an entirely new method.

In the next step, the room is replaced by a computer screen and the subject is replaced by a graphic representation of a turtle. This makes it possible for children to duplicate actions in their real world on the computer--and since children like to draw, the computer becomes nothing more than an easy-to-use drawing instrument.




The picture of the square from the practice room can be immediately created on the computer monitor by every member of the group:


```
forward 50
right 90
forward 50
right 90
forward 50
right 90
forward 50
```

These children were able to see the result of their work. Their learning was reinforced by the immediate feedback provided by turtle graphics and Logo. Logo is extremely interactive. Instead of displaying cryptic error codes or a brief error message, it tells you exactly what you did wrong in a friendly way.

Logo is an ideal example of an interpreted language because Logo gives confirmation of the execution of the instruction for each input. You'll be able to quickly acquaint yourself with this language once you have your own ST and Logo.



4.3.1 Procedures in LOGO

In Logo you can construct as many procedures as desired and thereby create a language tailored to your tastes, and application.

This process is not particularly complicated--you must think of an appropriate name for your new procedure and append *TO* to the procedure name. For example, the procedure **SQUARE** might look like this:

```
to square
forward 50
right 90
forward 50
right 90
forward 50
right 90
forward 50
end
```


To draw a square that is 50 units in length, you can tell the Logo interpreter to do this by typing the new procedure name: **SQUARE**.

But what if you want a square of a different size? You could have the procedure draw a square of any size by using variables. Here's the improved procedure:

```
to square :length
forward :length
right 90
forward :length
right 90
forward :length
right 90
forward :length
end
```

In this case, a variable named *length* is introduced. The **:** tells Logo that it is a variable. Within the body of the procedure, the occurrence of *:length* is replaced by the parameter that is tacked onto the procedure call. For example, **SQUARE 75** draws a square with a 75 unit side.

In the following pages, we'll take a brief look at the elements and predefined procedures in ST Logo.



4.3.1.1 LOGO Variables

Variables in Logo are designated by a prefixed quotation mark. Values are assigned to variables through the keyword `make`:

```
make "length 50
```

Logo recognizes both local and global variables. Local variables exist only within a procedure in which they are defined; once this procedure is exited, these variables are no longer accessible. The local variable `"length`, for example may be used to represent the number of characters of a word in one procedure and the dimensions of a square in a second procedure.


The instruction `local` makes an variable accessible to that procedure only:

```
>local "length
```

Another property of Logo is its ability of to view variables as executable statements. This is very helpful since it's possible for the computer to program itself.

```
make "angle [forward 50 right 90 forward 40]  
run :angle
```

The above example shows the method. A sequence of instructions is assigned to a variable which is then executed later with the `run` command.



4.3.1.2 Arithmetic operators in LOGO

All programming languages have some method of performing arithmetic operations like addition, subtraction, multiplication, and division. Dr. Logo supports the usual operators +, -, * and /.

The usual notation for arithmetic operations is like this: $2 + 3$

ST Logo also supports this notation: $+ 2 3$

Both yield a sum of five. The second notation is called *reverse polish notation*.

The functions `sin` and `cos` are available, where sine as well as cosine require as argument an angle given in degrees.

Other functions are:

`int`

returns the integer portion of a number by truncating it after the decimal point (the number is not rounded)

`random`

returns a random number between 0 and the value given as the parameter.

Example: `random 5` yields a number between 0 and 4.

4.3.1.3 Logical operators with LOGO

Logical operators are not often used by beginners, but are frequently used by more advanced programmers.

The operators `AND`, `NOT` and `OR` are available. The logical comparators `=`, `>` and `<` are also available.

With regard to logical operations, Logo resembles Pascal rather than BASIC. In BASIC, the true condition is represented as -1, while Logo and Pascal return `TRUE`.

4.3.1.4 Controlling the program execution

Programming in Logo is based on expanding the standard vocabulary. But Logo also requires commands for controlling the flow of a program's execution. In addition to the linear execution of program instructions, loops and branches are also possible. Here are the most often used:

bye

every Logo session should be ended with this command.

co

this command is an abbreviation for continue and causes a previously interrupted program to continue where it was interrupted

label name

identifies a program line by *name*. This label may then designate that line as the destination of a jump by *go name* statement.

go name

alters the flow of execution of the program; execution is continued at the line with the label identified as *name*.

**if expression command-list-1
command-list-2**

evaluates *expression* and executes *command-list-1* if the expression is true; otherwise perform *command-list-2* on following line.

op

interrupts the current procedure without changing the input value. This is then returned as the result of the procedure.

repeat n

performs the command list enclosed in parenthesis *n* times.
E.g. repeat 4 (fd 50 rt 90)

run command-list

executes *command-list* which follows.

stop

halts the execution of a program.

4.3.1.5 Graphics commands

With Logo, the user may divide the screen into different windows. The standard setting uses a variety of windows, one of which is set up as a graphics window and the others for communication with the user. The size of the windows can be changed, allowing the user to limit the field of movement of the turtle or to make the entire screen available for its movements. He can also eliminate the graphic screen and use the whole screen for text.

A few of the commands for manipulating the text screen:

ct

erases (clear text) the text window.

pr (*list*)

prints the individual elements of *list* to the screen. The outermost parentheses are removed and the individual list elements are printed followed by a carriage return.

setline

sets the line type to the input numbered style, width, color.

show (*list*)

similar to *pr*, but the outermost parentheses are not removed with *show*; *list* is displayed as the system sees it.

type

similar to *pr*, but the concluding carriage return is not used.

Here are some of the commands for manipulating the graphic screen:

cs

clears the graphic screen and sets the turtle to its starting position.

clean

clears (erases) the graphic screen however the turtle remains at its current position.

fence, window, wrap

One problem in working with computer graphics when plotting a mathematical function is that a point may lie outside of the drawing surface. Most computers respond with an error message unless the program is designed to *clip* or remove those points lying outside the addressable coordinates.

While it is possible to move the turtle outside the drawing surface, it can then no longer be seen. The advantage of this is that instructions like forward 5000 do not cause the procedure to crash.

Wrap causes the turtle to reappear on the left side of the screen if it disappears from the right side. A movement of forward 400 in the direction north (beyond the top screen edge) brings the turtle back to its starting position (in the 640x400 display mode).

Fence restricts the turtle's movement to the actual coordinates. If the turtle tries to go beyond the edge of the graphics window, the message "Turtle out of bounds" is displayed.

pal color-number

sets the color for the drawing pen to *color-number*.

setpal

selects a palette of pens from the enormous number of color shades which the Atari ST offers for drawing.

Here are some of the commands for controlling the turtle:

bk - back

fd - forward

these commands, followed by the length of the line in points, moves the turtle along this line in the current direction.

rt - right

lt - left

These two commands, followed by the angle of rotation in degrees, rotate the turtle at an angle in the appropriate direction.

ht - hide turtle

This command makes the turtle invisible during the drawing process. This speeds the drawing process since the computer no longer has to create and then erase the picture of turtle at each position.

st - show turtle

This command makes the turtle visible once again. Since all drawings are not made up of a single continuous line, the pen in Logo can be raised and lowered. The instructions necessary to do this are:

pu - pen up

pd - pen down

In addition, the pen can be replaced by an "eraser" in order to erase parts of the drawing. Do this with:

pe - pen erase

sets the drawing color to the background color so that the drawing can no longer be distinguished from the background.

In addition, ST Logo offers a drawing pen which reverses the color of the points.

px - pen reverse

automatically sets the unset points to the current pen color.

The turtle may be moved to a new position directly:

setpos - set position

this command, followed by the x and y coordinates positions the turtle to a specific location on the graphic screen.

The direction of the turtle can be reset:

seth - set heading

this command, followed by a direction turns the turtle to an absolute direction in degrees.

The status of the turtle can be examined:

tf - turtle facts

this command returns information about the turtle. This includes the turtle's coordinates, heading, pen status and pen color and indication as to whether the turtle is visible or not.

4.3.1.6. Primitives for word and list processing

ascii

returns the ASCII value of the first character of the following word.

char

returns the character corresponding to the following ASCII code.

bf - but first

returns all of the characters of the following word, except for the first character. For example, the function:

`bf "ATARI`

returns: **TARI**

bl - but last

returns all of the characters of the following word, except for the last character. For example, the function:

`bl "ATARI`

returns: **ATAR**

first

returns the first element of an expression. The element may be the first letter of a word or the first member of a list.

item *n*

returns the *n*th element of a list. For example, the function:

```
item 4 "ATARI
```

returns: **R**

count

returns the length of an expression. For example, the function:

```
count "ATARI
```

returns: **5**

If the expression is a list, the number of elements in the list would be returned:

```
count (monday tuesday wednesday thursday)
```

returns: **4**

empty *list*

returns either TRUE or FALSE depending on whether or not *list* is an empty list.

word *word*

returns either TRUE or FALSE depending on whether or not *word* is a word or a number.

word *list*

returns the concatenation of two or more words. For example, the function:

```
word "super "computer
```

returns: **supercomputer**

list (*elements*)

returns a list composed of *elements* including the parentheses.

sentence (*elements*)

returns a list composed of *elements* without the outermost parentheses.

Even with this incomplete list you can see that Logo offers a large number of built-in procedures and functions. In addition to all of the pre-defined keywords, **ST Logo** also offers instructions for controlling the synthesizer chip and for reading the joystick as well as the mouse input.

An additional command:

.contents


causes ST Logo to display its entire vocabulary including primitives and any extensions the programmer has made.

```

Desk File Run Edit Settings
LOGO DIALOGUE
?.CONTENTS
IDEGREES PENREVERSE SETPC PAUSE RIGHT ,DEF POREF LOADPIC WHERE => >= THING IF
ELSE <> <X HIDE TURTLE an empty word NODES QUOTIENT <= <= POALL WATCH RUN SUM
ETB6 BUTFIRST PPS POS NAMEP OUTPUT NOT ERALL ERASE PROCLIST HEADING TRACE EXP
CHANGEF PACKAGE LOCAL INT FALSE TOPLEVEL SIN ASCII BOX PIECE NOFORMAT ERN COS
UNBURY CATCH IFT TAN CLEAN LINEATTR EDALL LOG FILLATTR FENCE LABEL EMPTYP DIR
^ SETPOS PAL SETHEADING READLIST COPYON RERANDOM WINDOW ABS ARC SENTENCE IFF
ND FOLLOW PENERASE TURTLEFACTS EDF IFTRUE SETPEN SETPAN SORT SETPAL EQUALP SORT
T POTS STOP ,DEPOSIT LPUT TEXT SETY SETX POLY LOG10 TYPE POPS BURY SHOW REPEAT
RANDOM ERRACT PONS TEST TRUE FPUT ,REPTAIL DEFINEDP POTL > = YCOR LIST WORD <
XCOR POCALL PKGALL READCHAR ERPS WRAP ARCTAN KEYP ERNS REDEFP ,EXAMINE LAST
ETH CIRCLE SCREENFACTS SETZOOM SETTEXT ITEM SAVE / PATH - EDPS CLEARTEXT DEFIN
E + LEFT READQUOTE ,REPLACE EDNS *) HOME (UPPERCASE PX TT CLEARSCREEN ST FILL
EDIT RT LOWERCASE REMPROP PU TURTLETEXT GETTEXT TOWARDS RQ PR OR NAME PRODUCT
SHOWTURTLE LOAD LT BUTLAST PD OP CHAR MAKE RL ,PRM ,CONTENTS HT PENDOWN TF NUM
BERP SETPATH PI SF SE UC ER REMAINDER CT ,BUR COPYOFF 60 CS RC FORWARD PE ,AP
,FMT THROW NOWATCH PD SETLINE ,SPC SETFILL CO ,REM PPROP BACK ERASEFILE ,PKG
LC IF BL ELLIPSE PRINT SHUFFLE BK ,ENL LISTP WORDP PLIST NOTRACE SAVEPIC ERROR
COPYDEF FD ,PAK COUNT ED PRIMITIVEP MOUSE BF FIRST MEMBERP GPROP ROUND PENUP
RECYCLE GLIST RADIANS SETSCRUNCH POPKG

```


Since the entire working memory can be saved to floppy or on hard disk, you do not have to redefine the procedures at the start of each session. By building upon earlier procedures, you can create new enhanced Logo systems, such as a special games Logo, or a graphics Logo which might include procedures like BARCHART and FUNCTIONPLOT, or even an intelligent system which can carry on discussions with the user.



4.3.1.7 LOGO Summary

We hope that these pages have led you to the same conclusion that we found; that Logo is an excellent language for the beginner as well as for the advanced programmer.

A good programming language should be easy to learn. A few easy-to-remember expressions should suffice for simple tasks. You should not have to first master a large number of rules before defining variables.

The commands of the language should be easy to use. Its friendliness is enhanced if it is simple to use the peripheral devices such as the printer or disk drive.

A truly practical programming language should be easily expandable, so that extensions appear to be part of the standard commands. The user cannot distinguish between built-in and added commands.

A good language should be interactive so that the user can quickly develop and test new programs. In the event of program errors, it should display clear and helpful messages to the user.

All of these points speak for Logo. We can be happy that the Atari ST will offer both Logo and BASIC.

4.4 ST BASIC

Atari ST BASIC is one of the most feature packed BASIC languages available today. Not only is it somewhat compatible with GWBASIC (IBM), it also allows access to GEM, Graphics Environment Manager, and the VDI, Virtual Device Interface.

The manual that accompanies ST BASIC explains all of the commands, but is not for the novice programmer. If you have no previous experience programming in BASIC it is probably a good idea to purchase an introductory book covering ST BASIC. Experienced BASIC programmers will appreciate the large command set and its ease of use.

ST BASIC supports many of the capabilities of the ST computer including, windows, GOTO with labels, INPUT from various ports, PRINT USING, CHAIN, MERGE, RENUMBER, and many more advanced programming functions.

The command GEMSYS and VDISYS allow access to many of the graphic functions of the operating system. These commands in conjunction with the powerful TOS operating system make extremely difficult programming tasks very easy.


The speed of the 68000 processor and the fine implementation of BASIC will undoubtedly give rise to a vast amount of public domain software available free from your local user group or Atari dealer.

ST BASIC uses the GEM interface so the first time BASIC programmer should find it very easy to use. Experienced BASIC programmers will have to get used to the various windows that ST BASIC uses. The use of windows for OUTPUT, EDIT, COMMAND, and LIST should make learning the concepts required to program very easy to learn.

On the next page is a screen dump of the ST BASIC screen, showing its various windows and menus.

We'll not go into a very detailed description of ST BASIC in this book for there are many very excellent books available on the BASIC programming language. Abacus Software's *ST BASIC Training Guide* is an introduction to BASIC specifically written for the Atari ST.

Figure 4-1 ST BASIC Screen



We will now present a list of the commands available in ST BASIC to give you an idea of how flexible this language is.

ABS	ALL	AND
AS	ASC	ATN
AUTO	BASE	BLOAD
BREAK	BSAVE	CALL
CBDL	CHAIN	CHR\$
CINT	CIRCLE	CLEAR
CLEARW	CLOSE	CLOSEW
COLOR	COMMON	CONT
CONTROL	COS	CSNG
CVD	CVI	CVS
DATA	DEF FN	DEF SEG
DEFFBL	DEFINT	DEFSNG
DEFSTR	DELETE	DIM
DIR	ED	EDIT
ELLIPSE	ELSE	END
EOF	EQF	ERA
ERASE	ERL	ERR
ERROR	EXP	FIELD
FIELD#	FILL	FIX
FLOAT	FOLLOW	FOR
FRE	FULLW	GEMSYS
GB	GET	GET#
GO	GOSUB	GOTO
GOTOXY	HEX\$	IF
IMP	INKEY\$	INP
INPUT	INPUT#	INPUT\$
INSTR	INTIN	INTOUT
KILL	LEFT\$	LEN
LET	LINE INPUT	LINE INPUT#
LINEF	LIST	LLIST
LOAD	LOC	LOF
LOG10	LPOS	LPRINT
LSET	MERGE	MID\$
MKD\$	MKS\$	MOD
NAME	MEW	MEXT
NOT	OCT\$	OLD
ON	OPEN	OPENW
OPTION	OR	OUT
PCIRCLE	PEEK	PELLIPSE
POKE	POS	PRINT

PRINT#	PRINT USING	PTSIN
PTSOUT	PUT	QUIT
RANDOMIZE	READ	REM
RENUM	REPLACE	RESET
RESTORE	RESUME	RETURN
RIGHT\$	RND	RSET
RUN	SAVE	SGN
SIN	SOUND	SPACE\$
SPC	SQR	STEP
STOP	STR\$	STRING\$
SWAP	SYSDBG	SYSTAB
SYSTEM	TAB	TAN
THEN	TO	TRACE
TROFF	TRON	UNBREAK
UNFOLLOW	UNTRACE	USING
VAL	VARPTR	VDISYS
WAIT	WAVE	WEND
WHILE	WIDTH	WRITE
WRITE#	XOR	

Chapter 5

Introducing the 16-bit processors

- 5.1 The true 16-bit processors
- 5.2 The 68000 processor
- 5.3 The Instruction set of the 68000
- 5.4 Advanced features of the 68000

5. Introducing the 16-bit processors

This section will present a more detailed description of the actual components that make up the ST. These descriptions and comparisons should convince you of the amazing achievements of the Atari corporation in developing this amazing computer. Experienced programmers will discover that this new computer will be easy to learn and program.

Eight-bit microprocessors appeared on the market in the middle of the 70's. It wasn't long before these processors became the basis for the first home computers. Among the first were the PET 2001 (from Commodore), the Apple (from Apple Computer) both of which used the 6502 processors, and the TRS-80 (from Tandy) which used the Z-80 processor. These micros had from 4K to 16K of memory, a monochrome screen that displayed 25x40 or 16x64 characters and a built-in BASIC interpreter. They were priced at about one thousand dollars.

Now, for less than one-fourth of the price, you can get a computer with 64K of memory, built-in BASIC, high-resolution color graphics and excellent sound capabilities. This highlights the incredible improvement in performance and price that has taken place in such a short time. These were made possible mainly through advances in semiconductor technology.

For example, in the late 1970's increasing the memory capacity of an 8K PET to 32K cost well over \$300. Today, 64K of RAM costs less than \$25. The prices of peripherals have also fallen dramatically. In 1980, a printer could hardly be bought for less than \$750. Today, for about \$250, you can select from a wide variety of printers with far superior features and performance than was available only a few years earlier. This also holds true for mass storage devices. A dual disk drive with a 340K capacity cost \$1200 three years ago. Today, for about \$600, you can buy 2 megabytes of mass storage.

Why does it seem that we always need more and more storage capacity? In the past, mainframe computers typically had main memory capacities of from one-half to several megabytes. Using these computers, it was possible to execute computation intensive mathematical programs, usually written in the FORTRAN programming language, or to perform data intensive commercial tasks, most often written in the COBOL programming language. The compiler alone often required as much as 100K to run in such installations. If in 1975 you asked if FORTRAN programs could be

run on a home computer, the answer was definitely: "NO". But by 1980, FORTRAN and COBOL compilers were available for some microcomputers. The capabilities of these compilers was somewhat limited, but it was possible to use the huge base of existing FORTRAN and COBOL programs (these programming languages have been around since the fifties) on a microcomputer. But in so doing, the limitations of micros with limited memory also became clear.

The execution speed of compiled programs is considerably greater than interpreted languages such as BASIC. Therefore complex problem solutions which require a few minutes on a mainframe often take several hours on a micro.

Even when speed of execution is not a concern, the limited amount of memory - 64K - often prohibits the use of micros for some applications. To understand why 64K bytes represents the limiting boundary for most 8-bit microprocessors, we'll take a brief excursion into the hardware end of microprocessors.

When a microprocessor wants to access memory, it must determine which data it wants to read. To do this, each location in memory is referenced by a number or *address*. The processor has 16 address lines, each of which can assume one of two different conditions. Therefore you can select two different memory locations with each address line. If there are two lines available, then there are four different combinations for selecting or addressing the memory. With the sixteen address lines available, $2^{16} = 65536$ different memory locations can be addressed. Since $2^{10} = 1024$ addresses represent one kilobyte (K), 65536 is 64K. Since an 8-bit processor has eight data lines, each address can contain one of 256 (2^8) different values. A collection of eight bits is called a byte.

To go beyond these boundaries with an 8-bit processor requires some tricks. If you need more than 64K of memory for a program and/or data, then a section of memory not being used at that moment can be saved to a mass storage device, such as a disk, and reloaded later when required. The disadvantage of this procedure is that it is considerably slower than accessing the data directly in memory. The processor can access a memory location in mere microseconds while disk drive access (to retrieve the saved data) requires much more time and slows the execution of the program.

Another method of accessing more than 64K of memory is by *bank-switching*. Here, several memory banks each containing 64K are

used. Using a software-controlled switch, you can select between the different banks and thus gain access to more than 64K. This technique is certainly faster than the previous method, but complicates data access since the software must be specially written to switch between the memory banks.

For these reasons, people soon started looking for ways to improve the efficiency and capability of microprocessors. There were two major concerns: greater processing speed and the access of more memory.

A first step in the direction of greater memory access resulted in an improvement of the bank-switching technique. Switching between different memory banks was assigned to the processor instead of special software routines. The improvement uses a *segment register*. For example the 8086 and 8088 processors contain four segment registers. Using these registers, a 64K segment can be moved in main memory, which can amount to one million bytes (1M). Following each address, a segment register specifies the actual memory access: the code segment register determines from which segment the program instructions are to be fetched; the data segment register specifies the segment where the data is to be found. The stack segment register specifies the location of the stack and an extra segment register can point to an additional 64K segment. The disadvantage of this method is that only 64K can be accessed in one pass. The appropriate segment register must be reloaded in order to access memory outside of this range.

To increase the execution speed of the processor, the internal registers were widened from eight to sixteen bits. This allows operations, such as addition and subtraction, in the processor to be carried out on sixteen bits at a time. To retain this advantage during data transmission, the data bus was also widened from eight to sixteen bits in the case of the 8086. The 8088 is the low-cost version of the 8086, the data bus was retained at a width of eight bits. This permits a smaller size chip and reduces the number of other components, such as bus drivers, required on the board. The penalty is the reduced throughput, since sixteen-bit data must now be transferred with two sequential eight-bit transfers. Viewed in this manner, the 8088 no longer offers much advantages over an eight-bit processor, apart from an improved instruction set. This is why a BASIC program on the IBM PC runs only slightly faster than on many CP/M computers which have an 8-bit Z-80 processor.

Now Atari has changed all of this. The processor in the ST is a true 16 bit processor. This enables access to a greater amount of memory at extremely fast speeds. The ST stand for sixteen/thirty-two which accurately describe the 68000 processor.

5.1 The true 16-bit processors

Based on these considerations, the engineers at Motorola started the development of a sixteen-bit processor which would overcome these limitations. The goals were: a larger address range, greater data throughput (more speed), and a powerful instruction set which is both easy to learn and easy to use. The result of this development was the 68000, samples of which were already available as early as 1979. The 68000 is considered by many as the most capable of the sixteen-bit processors. It has already become the standard in many industrial applications. Because it forms the heart of the Atari ST, we'll examine the 68000 in greater detail in the next section.

Figure 5-1: 68000 Block diagram





Figure 5-2: 68000 Pin layout



5.2 The 68000 processor

In this section, we'll look closely at the makeup and operation of the 68000 processor. This may seem a bit technical but it will give you a better idea of the capability of the Atari ST. We'll also compare its operation to conventional eight-bit processors.

If you take a look at the circuit board in the Atari ST, you cannot overlook the 68000. It is the largest integrated circuit and is housed in a 64-pin dual-inline package. Why does the 68000 need so many "legs" (pins)? As previously mentioned, in order to transmit data at high speeds, the data must be sent in parallel, not in two sequential portions. Since the 68000 is a 16-bit processor, we need sixteen data lines. In order to address a large address range, the 68000 requires 24 address lines. This allows 2^{24} memory locations to be addressed. This is equivalent to 16,777,216 or 16 Mbytes, which is 256 times as much memory as a normal eight-bit processor can address. This huge address space is adequate for most any applications implemented on a microprocessor. As a comparison, a huge mainframe computer of the IBM System/370 can directly address 16 Mbytes.

So we find that 40 pins are needed just for the address and data lines. The remaining 24 pins are used for bus control, for the DMA access, for interrupt control, to output the processor status, and for peripheral control. A power source (only +5V) and the clock input are also required. The diagram in Figure 5-1 represents this graphically.

The effectiveness with which a processor can be programmed depends on the capability of the instruction set and on the number and type of internal *registers* available to the programmer. A register is a memory location within the processor which can be accessed at extremely fast speeds. It is in these registers that most of the operations which the processor can perform are carried out. This explains why the efficiency of a processor increases with the number and size of its registers.

To compare the 68000 processor to conventional eight-bit processors, Figure 5-3 shows the register set of the 6502 while Figure 5-4 shows the register set of the 68000. As you can see the 68000 has many more registers than the 6502. Additionally all of the 68000's registers are wider than 8 bits and can therefore process more data per operation.



Figure 5-3: 6502 Register set

Figure 5-4: 68000 Register set



As you can see, all of the registers in the 68000 are 32 bits wide. The register set is divided into eight data registers and eight address registers. Register A7 is available for two tasks and has a special significance. More about this later. The program counter is also 32 bits wide, although only the lowest 24 bits can be used. The program counter (PC) always points to the next instruction to be executed. The status register is sixteen bits wide and is divided into a *user byte* and a *system byte*.


As shown in the Figure 5-5 the data registers are all 32 bits wide too. But they can also be used as either sixteen or eight bit registers. If the contents of two data registers are to be added, for instance, you can specify whether the instruction will operate with eight, sixteen, or 32 bits.

When performing eight-bit processing, only the lowest eight bits are used; bits 8 through 31 remain unchanged and are not affected by the operation. If you use the 16-bit or *word* version of the instruction, only the lowest 16-bits are included in the operation and the upper 16 bits (16-31) are not affected. The 32-bit or *long-word* instruction uses the entire register.

These examples affect only two registers. To transfer data from a register to memory or from memory to a register, you can also specify the processing width. Since the processor has 16 data lines, transferring a single word (16-bits) is the easiest to accomplish. If only one byte (8-bits) are to be transferred, only data lines D0 to D7 are used. The processor informs the rest of the hardware via the bus control signals that only 8 bits will be transferred and that the contents of data lines D8-D15 is irrelevant. This means that it takes the same amount of time to transfer one byte of data as it does to transfer one *word* of data. But to transfer 32-bit data, it is no longer possible to transfer a single *long-word* in one move. First the processor transfers the most-significant word (bits 16-31); the address is automatically incremented; and then the least-significant word (bits 0-15) is transferred. The processor does this all automatically. To the user it looks as if he is working with a 32-bit processor. In this sense, the 68000 can be called a 32-bit processor by the same justification that the 8088 is said to be a 16-bit microprocessor.

The long-word transfer takes more time than the transfer of a 16-bit word, but it is faster than if you had to program the transfer of two words separately. Figure 5-5 clarifies the transfer methods.

Figure 5-5: 68000 Operands



5.3 The Instruction set of the 68000

To make full use of the 68000 hardware features, we need a capable instruction set. And to insure widespread use, this instruction set should be as easy as possible to learn.

These goals can be achieved with a relatively small number of basic instructions. Most of these instructions have slight variations to work with byte, word, or long-word data. The programmer needs to learn only a few instruction mnemonics. These mnemonics, may be appended by ".B" for byte, ".W" for word, or ".L" for long-word. The result is a simple, flexible instruction set that is easily learned. The operation codes for these instructions are each sixteen bits long. While this theoretically allows for up to 65536 (2^{16}) different instructions, this is clearly impractical. The following table lists the mnemonics for the 68000 instructions.

Table 5-1: The 68000 Instruction set:

ABCD	Add Decimal with Extend
ADD	Add
AND	Logical And
ASL	Arithmetic Shift Left
ASR	Arithmetic Shift Right
Bcc	Branch Conditionally
BCHG	Bit Test and Change
BCLR	Bit Test and Clear
BRA	Branch Always
BSET	Bit Test and Set
BSR	Branch to Subroutine
BTST	Bit Test
CHK	Check Register Against Bounds
CLR	Clear Operand
CMP	Compare
DBcc	Decrement and Branch Conditionally
DIVS	Signed Divide
DIVU	Unsigned Divide
EOR	Exclusive Or
EXG	Exchange Registers
EXT	Sign Extend
JMP	Jump
JSR	Jump to Subroutine
LEA	Load Effective Address
LINK	Link Stack (reserve stack area)
LSL	Logical Shift Left
LSR	Logical Shift Right
MOVE	Move Source to Destination
MULS	Signed Multiply
MULU	Unsigned Multiply
NBCD	Negate Decimal With Extend
NEG	Negate
NOP	No Operation
NOT	One's Complement
OR	Logical Or
PEA	Push Effective Address
RESET	Reset External Devices
ROL	Rotate Left without Extend
ROR	Rotate Right without Extend
RTE	Return from Exception

RTR	Return and Restore
RTS	Return from Subroutine
SBCD	Subtract Decimal with Extend
SCC	Set Conditional
STOP	Stop processor
SUB	Subtract
SWAP	Swap Data Register halves
TAS	Test and Set Operand
TRAP	Trap
TST	Test Byte
UNLK	Unlink (free stack area)

The 68000 recognizes the following conditions ("cc" in the instruction set):

The conditions marked with a * apply to two's complement arithmetic.

EQ	equal
NE	not equal
MI	negative
PL	positive
GT*	greater than
LT*	less than
GE*	greater than or equal
LE*	less than or equal
HI	higher than
LS	lower than or equal
CS	carry set
CC	carry clear
VS*	overflow
VC*	no overflow
T	always true
F	always false

So the 68000 instruction set contains no more mnemonics than a processor such as the 6502. This simplifies learning its machine language. The special capability of the 68000 lies in the numerous addressing methods as well as the power of the individual instructions. As you can see from the table, the 68000 has instructions for multiplication and division. These instructions are found in very few 8-bit processors. In order to briefly show you the different addressing modes, we'll illustrate them using what is probably the most universal instruction, **MOVE**.

The 68000 generally operates as a two-address machine. This means that a source (where the data comes from) and destination (where the data or result is placed) must be specified for each instruction.

The following instructions use the syntax of the 68000 assembler and demonstrate the different addressing modes.

002000 3200 MOVE D0,D1

This instruction copies the contents of register D0 to register D1. This instruction and those that follow, operate on a 16-bit word.

002002 323C0007 MOVE #\$0007,D1

This instruction loads the constant hexadecimal \$0007 into data register D1. As you see, this instruction is two words long. The opcode is \$323C, followed by the constant \$0007.

002006 7207 MOVEQ #7,D1

This instruction has the same effect as the previous one. The "Q" appended to the mnemonic stands for "quick." As you see, this instruction fits in only one word. It is limited to one-byte constants. The operand (07) is built into the opcode and makes the instruction shorter and faster to execute.

002008 3208 MOVE A0,D1

In this example, the contents of address register A0 are copied to data register D1. Data registers and address registers can be treated similarly.

00200A 32381000 MOVE \$1000,D1

In this example, the contents of memory location \$1000 are loaded into data register D1. Since the address can be represented as a 16-bit value, we speak of short addressing, similar to zero-page addressing on the 6502 processor.

00200E 31C11000 MOVE D1,\$1000

In this example, the contents of data register D1 are placed at address \$1000. This is, in contrast to the previous example, and transfers data from the processor to memory.

002012 3080 MOVE D0,(A0)

In this example, we are using indirect addressing. The contents of register D0 is stored at the memory location whose address is contained in address register A0. This again is a one-word instruction.

002014 31400032 MOVE D0,50(A0)

In this example, the addressing mode is similar to the previous one. The address at which the contents of D0 are stored is the sum of the contents of the address contained in register A0 plus the constant 50. You may think of A0 as containing the starting address of a table in which the 50th item is accessed.

002018 31801032 MOVE D0,50(A0,D1)

In this example, the addressing mode is further modified. The destination here is the sum of the contents of A0 and D1 and the constant 50.

00201C 33C000100000 MOVE D0,\$00100000

In this example, we return to the absolute addressing mode. This time the address cannot be represented as a 16-bit word, so an additional word is needed. This addressing mode is called absolute long.

002022 303A000A MOVE 10(PC),D0

Here, the program counter's relative addressing permits you to write relocatable programs. The address of the code to be executed is calculated from the current contents of the program counter plus the offset, here 10. The advantage of this mode is that no changes are necessary if the program is moved in memory; the tenth byte following the contents of the program counter is always accessed.

002026 30C0 MOVE D0, (A0) +

This addressing mode is very useful for working with tables. After the contents of D0 are stored at the address in A0, the contents of A0 are automatically incremented by 1, 2, or 4, depending on whether a byte, word, or long word was transferred. Afterwards, the next execution of the instruction automatically accesses the next table element.

00202A 32D8 MOVE (A0) +, (A1) +

This example shows indirect addressing with post-incrementation of both operands. The contents of the memory location contained at address register A0 is moved to the memory location contained in address register A1. Afterwards, both registers A0 and A1 are incremented. This allows you to transfer entire memory blocks simply. Note that the contents of a memory block is transferred directly to the destination address without having to be loaded into a register in the 68000.

00202C 48E7F7F8

MOVEM.L D0-D3/D5-D7/A0-A4, - (A7)

This instruction is unique to the 68000. Here the contents of registers D0-D3, D5-D7, and A0-A4 are placed on the stack with a single command. Address register A7 functions as the stack pointer. This instruction can replace up to 16 individual instructions for saving register contents on the stack, as is often required at the start of subroutines or interrupt service routines. The next instruction fetches the registers from the stack again.

002030 4CDF1FEF

MOVEM.L (A7) +, D0-D3/D5-D7/A0-A4

5.3.1 8-bit verses 16-bit

In order to give a concrete demonstration of the capability of this processor, let's perform the same task first with a 6502 processor and then with a 68000. As an example, we will copy a 4K block of memory from address \$3000 to \$4000.

```

; 6502 EXAMPLE
3000 SOURCE = $3000
4000 DEST   = $4000
1000 NUMBER = $1000 ; 4K BYTE
00F0 AD1    = $F0    ; POINTER TO SOURCE
00F2 AD2    = $F2    ; POINTER TO DEST
                *= $2000
2000 A9 00   LDA #<SOURCE
2002 85 F0   STA AD1
2004 A9 30   LDA #>SOURCE
2006 85 F1   STA AD1+1 ; INITIALIZATION
2008 A9 00   LDA #<DEST
200A 85 F2   STA AD2
200C A9 40   LDA #>DEST
200E 85 F3   STA AD2+1
2010 A2 10   LDX #NUMBER/256  LOOP COUNTER
2012 A0 00   LDY #0           ; INNER LOOP COUNTER
2014 B1 F0   LOOP LDA (AD1),Y ; SOURCE BYTE
2016 91 F2   STA (AD2),Y ; COPY TO DEST BYTE
2018 C8      INY
2019 D0 F9   BNE LOOP
201B E6 F1   INC AD1+1      ; INCREMENT POINTER
201D E6 F3   INC AD2+1
201F CA      DEX            ; OUTER LOOP
2020 D0 F2   BNE LOOP

```

Calculating the time that this 6502 program takes to execute, we find that it requires 65541 clock cycles. With a standard clock frequency of 1 MHz, this is 65.54 ms (milliseconds). Now let's "translate" the program into 68000 code.

003000	SOURCE	=	\$3000	
004000	DEST	=	\$4000	
001000	NUMBER	=	\$1000	
002000	207C00003000	MOVE.L	#SOURCE, A0	12
002006	227C00004000	MOVE.L	#DEST, A1	12
00200C	303C1000	MOVE.W	#NUMBER, D0	8
002010	12D8 LOOP	MOVE.B	(A0)+, (A1)+	12*4096
002012	51C8FFFC	DBRA	D0, LOOP	<u>10*4096</u>
		Total		90144

Even if you don't understand all of the details of the program, you can easily see that the 68000 program is shorter, in both the initialization and the actual program loop which performs the transfer. Since the 6502 only has an 8-bit register, a loop which is to be executed 4096 times is performed with two nested loops. Even loading the pointers with their starting values must be done in two steps on the 6502.

On the 68000, only one load command is required to initialize the registers for source, destination and number of iterations. The loop itself consists only of two instructions: the first copies a byte from the source to the destination address and automatically increments both pointers; the second decrements the loop counter and jumps back to the start of LOOP, as long as the counter D0 has not been decremented to zero.

The time to execute the individual instructions follows each instruction. We find that it takes 90144 clock cycles. With a 68000 driven at a standard clock frequency of 8 MHz, this takes 11.27 ms. This is already 5.8 times faster than the 6502. But we have not yet exhausted the power of the 68000. The above example transferred 4096 individual bytes.

But a 16-bit processor for such a task isn't very sensible. In the next example, we'll transfer 16-bit words in a single pass. This allows us to cut the number of loop iterations in half.

003000	SOURCE	=	\$3000	
004000	DEST	=	\$4000	
001000	NUMBER	=	\$1000	
002000	207C00003000	MOVE.L	#SOURCE, A0	12
002006	227C00004000	MOVE.L	#DEST, A1	12
00200C	303C0800	MOVE.W	#NUMBER/2, D0	8
002010	32D8	LOOP	MOVE.W (A0)+, (A1)+	12*2048
002012	51C8FFFC	DBRA	D0, LOOP	<u>10*2048</u>
			Total	45088

Now, we need only 45088 clock cycles or 5.64 ms to execute. This is about 11.6 times faster than the 6502. But the 68000 can work with long words, too. It naturally takes longer to transfer a long word than to transfer a word, but not as long as it takes to transfer two individual words. In addition, we can halve the number of loop iterations again. Here's the code:

003000	SOURCE	=	\$3000	
004000	DEST	=	\$4000	
001000	NUMBER	=	\$1000	
002000	41F83000	LEA.L	SOURCE, A0	8
002004	43F84000	LEA.L	DEST, A1	8
002008	303C0400	MOVE.W	#NUMBER/4, D0	8
00200C	22D8	LOOP	MOVE.L (A0)+, (A1)+	20*1024
00200E	51C8FFFC	DBRA	D0, LOOP	<u>10*1024</u>
			Total	30744

Now we need only 30744 clock cycles or 3.84 ms. This makes our 68000 program about 17.1 times faster than the 6502. In this example we used the LEA (Load Effective Address) instruction for initializing the address registers. This instruction is faster than the MOVE instruction and also requires fewer bytes. The processing width of the instructions in these examples were indicated by appending ".B", ".W", ".L".

In general, we can say that the 68000 is about 10 to 20 times faster in program execution than most popular 8-bit processors. In addition, the programs are usually shorter and easier to read thanks to the powerful and highly orthogonal instruction set. In special cases, the speed advantage over 8-bit processors can be even larger.

5.4 Advanced features of the 68000

The 68000 can be operated in two modes. The first is the *user mode* and the second is the system or *supervisor mode*.

This permits a strict separation between the operating system and user programs. Some instructions are not allowed in the user mode--they are *privileged*. If an attempt is made to execute these instructions anyway, the 68000 branches to an exception handling routine similar to an interrupt. This routine, which is part of the operating system, can then react accordingly, by outputting a message, for instance. Privileged instructions include the STOP command which halts the processor, and the RESET command which resets the peripheral devices.

In general, these privileged instructions are reserved for operating system functions that protect important computer resources. For example, if valuable data is kept in a computer system, you may want to secure it from being read. If all of the facilities of the operating system are available to a user, then he may be able to overcome the security. But if some instructions are not available to him, then it is possible to keep him from accessing the data. This is one of the main purposes of privileged instructions.

As you can see from the register assignments, there is a different stack pointer, A7 for user and supervisor. They are referred to as the USP (user stack pointer) and SSP (system stack pointer). The operating system stack pointer cannot be changed by a user program, only from supervisory mode.

Similarly, interrupt priorities can be changed only from the supervisory mode. The 68000 has seven interrupt priorities. If the interrupt priority is zero, then all interrupts are permitted. Under the highest interrupt priority, only the non-maskable interrupt is permitted. If the interrupt mask is set to 3, for example, only interrupts 4 to 7 are permitted. If an interrupt occurs, the priority is automatically set to the value of this interruption during the execution. In this way, an active routine servicing an interrupt cannot be interrupted by the occurrence of one with lower or equal value. The processor is automatically switched to supervisory mode when an interrupt occurs.

Another feature of the 68000 is the `CHK` instruction. This instruction continuously checks the contents of an address register and if it falls outside of a desired address range, causes an interrupt. This is used in multi-user applications for example, to make sure that each user accesses only his assigned area of memory. If not, the `CHK` instruction can branch to an operating system routine which can then react accordingly.

Another use for this instruction is for high-level language compilers. Here, the instruction can be used to check if the index of an array remains within the declared limits and to output an error message if it does not.

The 68000 has a hardware feature which branches to an exception routine if an attempt is made to access an address range not installed on the computer. Since the 68000 can indicate if it is in the supervisor mode or user mode via the status lines, you can create a simple circuit if the processor attempts to address a certain area in user mode. This is done in the Atari ST. So it's not possible to access the stack area of the supervisor or the peripherals from the user mode. If an attempt is made, the bus error routine of the operating system is activated.

The 68000 has features to support high-level programming languages in which recursive calls are used. The `LINK` instruction can be used to reserve a data area to save the subroutine's local variables on the stack. The stack area can be freed again after subroutine is finished with `UNLK` instruction.

5.4.1 Relocatable programs

Since systems with a few hundred kilobytes of memory such as the Atari ST generally use only disk operating systems (DOS), the user programs are loaded into RAM. If several programs or program modules are combined, it can no longer be guaranteed that programs are always loaded at the address for which they were originally written. To do this one needs either a linker--a program which recalculates the absolute addresses of a program for another memory area--or a program which can run anywhere in memory. Such a program is said to be *relocatable*. The 68000 is ideally suited for writing such programs because it has program-counter relative instructions for subroutine calls, branches, and data accesses.

In conclusion, we can say that the 68000 is excellently suited for larger microcomputer systems with lots of memory. It supports the demands which are placed on such a system today. It offers separation between operating system and user programs. So it is well-suited for multi-user applications as well as for the implementation of modern high-level languages.

Chapter 6

The Architecture of the ST

- 6.1 ST Interfaces
 - 6.1.2 Floppy disk interface
 - 6.1.3 The Hard Disk
 - 6.1.4 High speed through DMA
 - 6.1.5 The Printer Interface
 - 6.1.6 The Serial Interface
 - 6.1.7 The MFP 68901
- 6.2 The sound capabilities of the ST
 - 6.2.1 The MIDI Interface
- 6.3 The ST Keyboard
 - 6.3.1 The ST Keyboard processor
- 6.4 The ST Video Interface

6. The Architecture of the ST

Now that we have learned a few fundamental things about the 68000 CPU, let's look at the layout of the ST in more detail.

Let's begin with the ST's memory layout. As you already know, the 68000 can address 16 MB of memory. In the 520 ST, 512K or one-half megabyte of RAM is available. This is only 1/ 32nd of its available address range. The RAM area starts at address 0 and goes to address \$07FFFF or decimal 524,287. The 1040 ST has 1Megabyte of RAM which ranges from address 0 to \$09FFFF or decimal 1,023,999. One-half megabyte is eight times as much memory as an 8-bit computer with 64K. On the ST, almost all of this memory is available to the user, as we will see later. On other 16-bit computers with 128K of RAM, more than half of the memory is often taken up by the operating system so that often no more than 32K is available for the user.

Not so with the ST. Here, almost all of the operating system is in ROM (early versions of the ST loaded the operating system from disk.) This ROM area is large, 192K, and lies from address \$FC0000 to \$FEFFFF (decimal 16,515,072 to 16,711,679). This ROM contains the BIOS, the actual operating system TOS and GEM (more about this later). Only the scratch-pad memory needed by the operating system and the screen storage lie in RAM.

The fact that the operating system is built into ROM has several advantages. For one, the operating system does not have to be loaded from disk. For another, it does not occupy any RAM, which results in more available RAM for the user. With an address range of 16M, the space consumed by the ROM can be ignored since it is just one percent of the total address range allowable.

For those not satisfied with this, the address range from \$FA0000 to \$FBFFFF (decimal 16,384,000 to 16,515,071) is reserved for ROM cartridges. User programs as well as operating system expansions such as additional languages can be placed in this 128K range. The advantage over loading from disk is again clear: the loading procedure is avoided and no RAM is consumed.

Figure 6-1: ST memory map

ATARI ST Memory Map

\$FF FC00	I/O - Area	16776192
\$FF FA00		16775680
\$FF 8800	I/O - Area	16746496
8600		16745984
8400		16745472
8200		16744960
\$FF 8000		16744448
\$FE FFFF	192 K System ROM	16711679
\$FC 0000		16515072
	128 K ROM Expansion Cartridge	
\$FA 0000		16384000
\$09 FFFF	1Mb RAM 512 K RAM	1023999 ATARI 1040 ST
		524287 ATARI 520ST
\$07 FFFF		
\$00 0000		0

The range from address \$FF8000 (decimal 16,744,448) to the end of the physical address range is reserved for peripherals. There are 512 bytes set aside for each device. This is a total of 32K for peripherals, 0.2% of the total address range. The Figure 6-2 shows how the I/O area is divided.

Of the 1M or 512K (depending on the model), only 32K is required for the screen RAM. This RAM contains the bit map which serves for the representation of the picture on the monitor. The screen area can lie anywhere in RAM, thereby making it possible to switch between several screens, for instance.

Figure 6-2: I/O Assignments

\$FFFC00	2 ACIA's 6580
\$FFFA00	MFP 68901
\$FF8800	SOUND AY-3-8910
\$FF8600	DMA / WD 1770
\$FF8400	RESERVED
\$FF8200	VIDEO CONTROLLER
\$FF8000	DATA CONFIGURATION


6.1 ST Interfaces

On the next page is a block diagram of the ST. Using this as a starting point, we'll examine the peripheral interfaces and devices.

As an overview, we first list the major interfaces:

- * TWO 3 1/2" DISK DRIVES
- * DMA INTERFACE FOR CONNECTION OF A HARD DISK WITH 10 MB STORAGE CAPACITY
- * CENTRONICS PARALLEL INTERFACE FOR CONNECTION OF A PRINTER
- * SERIAL RS 232 INTERFACE FOR CONNECTION OF DEVICES SUCH AS A MODEM OR A PRINTER
- * MIDI INTERFACE FOR CONTROL OF EXTERNAL MUSIC SYNTHESIZERS
- * 2 JOYSTICK PORTS, ONE FOR CONNECTION OF A MOUSE
- * VIDEO CONNECTIONS FOR RF (TELEVISION), B/W MONITOR, AND RGB FOR COLOR MONITOR
- * CONNECTION OF AN INTELLIGENT KEYBOARD

Figure 6-3: Block diagram of the ST



6.1.2 Floppy disk interface

The most important peripheral device for a microcomputer is the mass storage device. The ST allows up to two floppy disk drives to be connected, both of which use 3 1/2" diskettes. Disk drives of 360K and 720K are available. The ST uses the more compact 3 1/2" disk instead of the more common 5 1/4" disks. They offer several advantages. Even though they have about the same storage capacity, they are smaller than the 5 1/4" floppies. The diskettes are also less sensitive to mechanical damage since they are contained in a hard plastic case and so the magnetic surface is better protected. The opening for the read/write head on the diskette is also protected by a metal shutter which moves aside when the disk is inserted in the drive. In place of a write protect notch of the minifloppies, the 3 1/2" diskettes have a slide to protect the diskette from being written to. The smaller drives allow a smaller physical package and reduced cost.


The ST has a built-in controller for operating the disk drives (WD 1770 from Western Digital). No special interface is necessary to connect the drives--a simple cable suffices. The disk controller receives its commands from the processor. The disk controller does not get the data to be written to and later read back in from the processor, but from a DMA component. The next section about the hard disk will further explain the DMA.

The disk controller is responsible for the way in which the individual bits are written to the diskette and for reading them in again later. The processor gives the disk controller commands such as "read sector" or "write sector". The disk controller can also format diskettes. The processor supplies the data which specifies the desired format of the diskette: how many tracks are used; how many sectors per track; how many bytes per sector; etc. The WD1772 uses the System/34 format and the disk format is very similar to the MS-DOS disks on the DATA GENERAL/ONE portable. You can also modifying the disk drive hardware and install an IBM 5 1/4" disk drive. The IBM type drive can then be used to transfer data from the two computers but you will not be able to format IBM disks with the ST.

In addition, the disk controller is responsible for data integrity. To detect read errors, the disk controller creates a checksum called CRC (Cyclic Redundance Check) bytes, which are written to the disk following the data. Later, these CRC bytes are compared with the value calculated when the data is reread. If these two values are not identical, then the disk controller tells the processor via its status register that an error occurred.

The floppy disk has been the standard mass storage device for a microcomputer, and can be used for storing programs and data.

Figure 6-4 3 1/2" Floppy Disk



6.1.3 The Hard Disk

If you work with a 520 ST and a 360K disk drive, you'll know that the diskette has roughly the same capacity as the RAM capacity of the computer. If you use the computer for word processing, or you want to save entire graphic pages of 32K each on diskette, you quickly reach the storage capacity of the diskette. The hard disk is a way to overcome this limitation. Atari offers a hard disk for the ST with 10 MB capacity.

What are the major differences between a hard disk and a floppy?

The most spectacular difference for the user is the enormous storage capacity. The 10 MB hard disk offered for the ST holds 20 to 40 times as much data as will fit onto a floppy. How is this possible?

In contrast to a floppy, a hard disk consists of a hard, rigid platter, usually made of aluminum and covered with a thin magnetic layer. The diameter of the platters on the Atari drives is the same as the floppys--3 1/2". The data can be written to the rotating platter or read back with a magnetic head, as with a floppy drive.

The essential difference between the hard disk and a floppy is that the hard platter rotates at 3000 RPM instead of 300 revolutions per minute.

Another difference is that the read/write head in the hard drive never touches the platter but "flies" over it at the very low altitude of half a micrometer (0.0005 mm). Such a method naturally requires a great deal of mechanical precision. A dust or smoke particle that comes between the head and the platter can damage the platter and result in subsequent data loss. For this reason, the platter is enclosed in a hermetically sealed case and cannot be changed like a normal diskette.

Coupled with higher quality components and increased speed of rotation, the track density of a hard disk is more than 800 tpi (tracks per inch) while the recording bit density is over 8000 bits per inch. The corresponding track density for a 3 1/2" floppy is typically 135 tpi.

As a result of the denser track layout, the hard disk also reduces the time the head needs to move from one track to another. This results in faster access time. A more important factor in performance is the data transfer rate. Because the platter rotates ten times faster than the floppy and is recorded at a higher density, the hard disk can deliver the data with a speed of about one

megabyte per second. With a floppy, the data can be read at a speed of "only" about 25K per second.

In order to transfer data from a storage medium to the computer, or the other way around, peripheral interfaces are normally used to execute these tasks. To read external data, the processor reads the contents of the ports of the interface and then writes the value into memory. The I/O interface informs the processor via special registers when the next byte is ready and can be fetched by the processor. This method has proven itself and works quite well. Depending on the processor, transfer rates up to 100K per second can be obtained.

6.1.4 High speed through DMA


In order to make full use of the transfer speed another method has been devised. The method is called "Direct Memory Access" or DMA for short. What's this all about? Our goal is to transfer data from mass storage to the memory of the computer. A way has been found to do this directly without involving the processor. There is a special component for this task, the DMA controller. This device has the job of bypassing the processor and writing the data coming from the hard disk directly into the memory of the computer. The DMA controller can also do the reverse when writing to the hard disk--the data is read from memory and sent directly to the hard disk.

The DMA controller is a custom I/O component for the processor. To send data to the hard disk, for instance, the processor must program the DMA controller appropriately. Among other things, the DMA controller must know the address range which is to be sent to the hard disk. When it has received the appropriate command, the DMA controller informs the processor via certain control lines that it now wants to access the data and address lines. The processor frees the bus and goes into a wait state. Now the DMA controller can access the entire memory range, read the data and send it to the hard disk. This occurs at the same high speed which the processor reads data from the memory. When the transfer is done, the DMA controller frees the bus and the processor can continue. On the ST, a data transmission rate to and from the hard disk of 1.33 MB per second is possible. This allows the entire contents of the memory to be send to the hard disk in less than half a second!

For the sake of fairness we want to explain the disadvantages of a hard disk. For one, we cannot change the platter. This makes a hard disk unsuited for exchanging data between similar computers. The second and more important limitation involves the security of the data. If the hard disk should ever become defective, more than 10 MB of data may be destroyed. This can happen if a dust particle gets between the head and the platter, for instance. When something like this happens it is called a "head crash". In connection with the data integrity issue, there is often no simple way of backing the entire hard disk up. There are tape drives sold for this purpose, but they are frequently more expensive than the hard disk.

Since you will usually work with one floppy and one hard disk, your most important data should be copied from the hard disk to floppy diskettes at regular intervals. Hard disks today have quite a high quality standard and errors rarely occur.

Figure 6-5 DMA Port




6.1.5 The Printer Interface

So that you can print out your work, in black and white or color, the ST has a printer interface. With a printer you can reproduce one or as many documents as necessary. More interesting though is the possibility to output graphics as hardcopy. Because (as you will see) the screen of the ST is always represented as graphics, the resolution in monochrome mode is 640 points horizontally by 400 points vertically, it is possible to output the contents of any desired screen to paper with a dot-matrix printer capable of graphics.

To connect a printer to a computer, both must have compatible interfaces. The Centronics parallel interface is considered the standard interface which printers and computers use. The Centronics interface is a parallel interface in which the data is transmitted byte by byte. Each character, such as a letter or digit, is represented in the computer with something called an ASCII code. In this code, each character is denoted through one byte or 8 bits. This makes it possible to represent a total of 256 different characters. For example, the letter A is represented by the ASCII code 65. When a character is to be transferred from the computer to the printer, the printer must be sent the corresponding byte of the ASCII code. With the Centronics interface, each bit has its own line. Two handshake lines are used so that the computer and printer can agree on the time of the transfer.

If the computer wants to send a character to the printer, it places the data on the interface lines and puts a low signal on the STROBE line for a brief time (about one microsecond). This tells the printer that a character is ready and can be read from the interface. The BUSY line is required so that the computer knows when the printer has accepted the byte so it can send the next byte. As long as the printer is busy with the accepting and processing of a character, it makes the BUSY line high. The computer need only wait until the BUSY line goes back to low before it sends the next byte with a STROBE pulse.

On older operating systems installing printers could be a difficult and time consuming task. Atari had the good sense to include an Epson {and compatibles} printer driver in its operating system. The printer can be installed from the desktop, here is the installation menu.



Here are the lines and the data directions.

COMPUTER	PRINTER
PSG PORT A6	STROBE ==>
PSG PORT B0	D0 ==>
PSG PORT B1	D1 ==>
PSG PORT B2	D2 ==>
PSG PORT B3	D3 ==>
PSG PORT B4	D4 ==>
PSG PORT B5	D5 ==>
PSG PORT B6	D6 ==>
PSG PORT B7	D7 ==>
MFP IO <==	BUSY
GND	

PSG=Programmable Sound Genarator
MFP=Multi-Function Peripheral


6.1.6 The Serial Interface

In contrast to parallel data transfer, which has a line available for each bit, only one line is required for the data with serial transmission. When a byte is to be transferred, the bits are transmitted one after the other over the same line. When sending such a bit stream, the receiver must know with which bit the transmission of a new byte starts.

Two procedures have been developed to solve this problem. The first is called synchronous transmission. Here the sender and receiver are supplied with the same clock so that they, as the name implies, operate synchronously and can assign the bits received to the bytes properly. Since this method requires a coupling between sender and receiver, it is seldom used in microcomputers.

Much more widely used is the asynchronous method. Here the sender and receiver need only be connected to each other over one data line and a common ground line. This data line is at a specific condition in the wait state. If the sender wants to send a byte, it first sends a start bit by placing the data line low. Now the receiver knows that a data byte follows and can prepare to receive it. The individual bits are now transmitted over the data lines with low or high signals depending on their value. After the transfer the sender sends one or two stop bits, which at the same time correspond to the rest state. When the next byte is to be transferred the whole procedure starts over from the beginning. The figure below should clarify this procedure.

Serial transmission



With this method any amount of time may pass between the individual bytes. A prerequisite of this procedure is that the receiver must constantly watch the data line so as not to miss the start bit. After a start bit is received it "collects" the data bits and puts them together into bytes. In order to inform the processor that a data byte has arrived, it generates an interrupt which causes the processor to get the complete byte from the peripheral component and write it to a buffer for later processing, for example. But let us take another look at the actual transmission.

The individual bits after the start bit have a set length. The sender and receiver must agree on this length before the transmission is attempted because the receiver has no way of determining this length. In our example, 8 bits are used. Specific values have been established for the times of each bit. The shorter the time for each bit, the greater the transmission speed. The speed is measured in bits per second. The name baud has been established for this.

BAUD RATE	BIT LENGTH (MS)
500	20.00
110	9.09
150	6.67
300	3.33
600	1.67
1200	0.83
1800	0.56
2400	0.42
3600	0.28
4800	0.21
7200	0.14
9600	0.10
19200	0.05

These are the baud rates at which the ST can send data. While the Centronics printer interface is designed only for sending information, the serial interface is suited for receiving data as well. A second data line is used for this which the ST uses as a serial data input. This input can be programmed for the same transmission speeds.

This serial transfer method is usually called the RS-232 standard. It functions faultlessly as long as the receiver processes the data as quickly as the sender transmits it. If the receiver is busy processing earlier data when the next byte arrives, the incoming data is lost without the sender even being

aware. To prevent this loss of data several techniques termed handshaking are used. One method employs hardware and two additional lines, the "Clear To Send" or CTS line and "Request To Send" or RTS line. When the receiver is ready to receive a byte, it signals this through an appropriate signal on the RTS line. The sender first checks the line before making a transmission and waits as long as required until the receiver can accept the next byte. This procedure is often used when connecting terminals or printers. The disadvantage is that additional lines are required.


The Atari ST supports the following handshake lines for the RS-232 interface:

PIN	ABBR.	SIGNIFICANCE	I/O COMPONENT
4	RTS	REQUEST TO SEND	PSG PORT A3
5	CTS	CLEAR TO SEND	MFP I2
8	DCD	DATA CARRIER DETECT	MFP I1
20	DTR	DATA TERMINAL READY	PSG PORT A4
22	RI	RING INDICATOR	MFP I6

Another method which has been gaining more and more use is data transmission over the telephone. Here the individual bits are converted into different tones, transmitted over the phone lines, and converted back into voltage signals by the receiver. Devices which make such a conversion possible are called modems. When using the telephone, however, we have only one line running in each direction. The hardware method of handshaking described above is therefore impossible. A software method has been devised that allows handshaking over the phone lines. If the receiver wants to halt the transmission, it sends the character XOFF. This character has the ASCII code 19 and can be entered as a Control-S. Once the receiver is ready again to accept more data, it lets the sender know this by sending it XON which has the ASCII code 17 and can be entered as Control-Q.

The Atari ST supports the RS-232 interface through the operating system. The baud rate, number of data bits to be transmitted, and the type of handshaking can be programmed through the appropriate calls or selected from the desktop. You choose between no handshake, RTS/CTS, or XON/XOFF protocol. In addition it is possible to check the bytes transmitted with a parity bit.

Again the TOS operating system allows these parameters to be changed easily from the desktop.



6.1.7 The MFP 68901

So far in this chapter we have talked about the peripheral components which are responsible for transferring data between the computer and the outside world. As an example of such chips, we want to take a closer look at the MFP 68901.

MFP is an abbreviation for "Multi-Function Peripheral". The name implies that this chip can perform several tasks. The MFP 68901 is a relatively new I/O component in the 68000 family.

These are the chip specifications:

- * 8 individually programmable input/output lines with interrupt capability**
- * Interrupt controller for 16 interrupt sources with individual masking**
- * Four timers, two of which have multiple functions**
- * The timers can be used as baud-rate generators for the serial channel**
- * One channel synchronous and asynchronous serial input/output (Universal Synchronous/Asynchronous Receiver/Transmitter, USART)**

The registers of a peripheral component are addressed by the processor in exactly the same way as memory locations. The MFP 68901 has 24 internal registers, the significance of each is shown in TABLE 6-1:

Table 6-1: MFP 68901 Internal Registers

REGISTER	ABBR.	DESIGNATION
1	GPIP	GENERAL PURPOSE I/O REGISTER
3	AER	ACTIVE EDGE REGISTER
5	DDR	DATA DIRECTION REGISTER
7	IERA	INTERRUPT ENABLE REGISTER A
9	IERB	INTERRUPT ENABLE REGISTER B
11	IPRA	INTERRUPT PENDING REGISTER A
13	IPRB	INTERRUPT PENDING REGISTER B
15	ISRA	INTERRUPT IN-SERVICE REGISTER A
17	ISRB	INTERRUPT IN-SERVICE REGISTER B
19	IMRA	INTERRUPT MASK REGISTER A
21	IMRB	INTERRUPT MASK REGISTER B
23	VR	VECTOR REGISTER
25	TACR	TIMER A CONTROL REGISTER
27	TBCR	TIMER B CONTROL REGISTER
29	TCDCR	TIMER C AND D CONTROL REGISTER
31	TADR	TIMER A DATA REGISTER
33	TBDR	TIMER B DATA REGISTER
35	TCDR	TIMER C DATA REGISTER
37	TDDR	TIMER D DATA REGISTER
39	SCR	SYNCHRONOUS CHAR. REGISTER
41	UCR	USART CONTROL REGISTER
43	RSR	RECEIVER STATUS REGISTER
45	TSR	TRANSMITTER STATUS REGISTER
47	UDR	USART DATA REGISTER

Since the MFP has an 8-bit data bus, the individual registers are numbered by two's. With the help of these registers, the processor can control the operating mode of the MFP, supply it with data to be output, and get input data from it.

A peculiarity of the MFP 68901 is the 16-channel interrupt controller. It can manage 16 interrupt sources and permit or mask (prohibit) them individually. These interrupts are prioritized, meaning that the interrupt sources have varying weights. The MFP 68901 can generate a unique vector number for each of these interrupts, so a separate interrupt service routine is available in the 68000 for each of the interrupts. The interrupt priorities are all the same as far as the 68000 is concerned; the prioritizing is taken care of in the 68901. The 68901 uses the highest level of maskable interrupt in the 68000, namely six.

The vertical sync interrupt resides at interrupt level four and the horizontal sync interrupt is at level two. These interrupts work as auto-vector interrupts which means that the vector number is determined by the 68000. The Atari ST uses only the top two of the three interrupt-level inputs, resulting in the fact that only levels 2, 4, and 6 can be used. Level 7, with the highest priority represents the NMI (non-maskable interrupt).

The MFP 68901 can manage 8 internal and 8 external interrupt sources which have the following priority:

Table 6-2: MFP 68901 Interrupt priority

PRIORITY	CHANNEL	DESCRIPTION
HIGHEST	15	EXTERNAL INTERRUPT 7 I7
	14	EXTERNAL INTERRUPT 6 I6
	13	TIMER A
	12	RECEIVER BUFFER FULL
	11	RECEIVER ERROR
	10	SEND BUFFER EMPTY
	9	TRANSMISSION ERROR
	8	TIMER B
	7	EXTERNAL INTERRUPT 5 I5
	6	EXTERNAL INTERRUPT 4 I4
	5	TIMER C
	4	TIMER D
	3	EXTERNAL INTERRUPT 3 I3
	2	EXTERNAL INTERRUPT 2 I2
	1	EXTERNAL INTERRUPT 1 I1
LOWEST	0	EXTERNAL INTERRUPT 0 I0

Each of these interrupt sources can be enabled or disabled by programming the MFP 68901. The eight external interrupt sources use the data lines of the 8-bit port as inputs. If all of the lines are not used as interrupt inputs, the rest can be used as normal port lines. Internal interrupt sources can be the timers A through D, which can be used in different operating modes. They can create delay times, measure pulse widths, generate square waves, or count events.

In addition, two of the timers can be used as baud rate generators for serial transmission. The 68901 can generate an interrupt when an entire byte has arrived serially. The byte can then be input from the USART data register by an appropriate routine. The MFP 68901 can call an interrupt routine, using interrupt channel 10, which places the next byte at its disposal when it has output a complete byte.

The MFP 68901 is housed in a 48-pin case and is connected to the asynchronous bus of the 68000. The timers can be supplied with a clock independent of the clock frequency of the processor. In the ST the 68901 serves as the serial interface and works as an interrupt controller. The 16 interrupt levels of the 68901 are assigned as follows:

Table 6-3: Interrupt structure of the ST

PRIORITY	CHANNEL	SIGNIFICANCE
HIGHEST	15	I7 MONOCHROME MONITOR DETECT
	14	I6 RS-232 RING INDICATOR
	13	TIMER A, SYSTEM CLOCK
	12	RS 232 RECEIVER BUFFER FULL
	11	RS 232 RECEIVER ERROR
	10	RS 232 SENDER PUFFER EMPTY
	9	RS 232 TRANSMISSION ERROR
	8	TIMER B, HORIZONTAL LINE RETURN
	7	I5 FLOPPY DISK CONTROLLER
	6	I4 ACIA 6850, KEYBOARD AND MIDI
	5	TIMER C
	4	TIMER D, RS 232 BAUD RATE GENERATOR
	3	I3 GPU OPERATION DONE
	2	I2 RS 232 CLEAR TO SEND
	1	I1 RS 232 DATA CARRIER DETECT
LOWEST	0	I0 CENTRONICS BUSY

6.2 The sound capabilities of the ST

The ST contains a built-in sound generator of type AY-3-8910 from General Instruments or a replacement type YM 2149 from Yamaha.

The 40-pin chip has the following specifications:

- * three independently programmable tone generators
- * a programmable noise generator
- * outputs which are completely software controlled
- * programmable mixer for tone and noise
- * 15 logarithmic volume levels
- * programmable waveforms
- * two bi-directional 8-bit ports

The programmable sound generator AY-3-8910, abbreviated to PSG, is constructed as follows:

Three tone generators A, B, and C create the square waves for each channel. A noise generator produces a square wave with a random pulse width. Three mixers can mix the outputs of the tone generators with the noise generator. The D/A converter can be controlled with either the amplitude control or the waveform generator. The programmable waveform generator can modulate the tone outputs into different wave forms. The D/A converter creates 16 different volumes determined by the amplitude control.

The PSG has 16 registers which can be programmed by the processor:

Registers R0 through R5 serve to determine the pitch of the three channels A, B, and C. Registers R0 and R1 control channel A, R2 and R3 are responsible for channel B, and R4 and R5 for channel C. The value for the period of the tone is written to these registers as a multiple of the clock frequency of the PSG divided by PSG. Since the PSG in the Atari ST is driven at 2 MHz, the basic value for the period duration is 8 microseconds. Only twelve bits of

the 16 bits in the two registers can be used; the top four bits are ignored. What frequencies can be created? We can write values between 1 and 4095 to the registers and we get the following values:

$$\begin{array}{rcl} 1 * 8 = 8 & \text{microseconds} => & 125,000 \text{ Hz} \\ 4095 * 8 = 32760 & \text{microseconds} => & 30 \text{ Hz} \end{array}$$

We can therefore create frequencies which far exceed the audible range.

Register R6 sets the basic frequency of the noise generator. Here only the 5 lowest bits are used.

Register R7 selects the signals of the tone generators and noise generator. With six bits you can specify which channels of the tone generator are to be switched to the corresponding outputs and to which channels the noise generator are to be added. Bits 6 and 7 are responsible for the data direction of the two 8-bit ports.

Registers 8 thru 10 are the amplitude control.

Registers 11 and 12 are used for programming the waveform.

The remaining registers are used for the I/O ports. The 8-bit port B of the PSG (register 15) is used in the Atari ST for the Centronics interface. Port A (register 13), which can only be used as output, delivers the handshake signals for the serial and parallel interfaces. Three bits of port A are also used for the floppy interface to select the drives and heads.

ADSR control is possible with the programmable waveform. This is an abbreviation for Attack-Decay-Sustain-Release. This allows the attack time, decay time, sustain time, and release time of a tone to be set.

The ST operating system contains procedures for passing the appropriate parameters to the PSG.

6.2.1 The MIDI Interface

Another feature of the ST is the built-in MIDI interface. MIDI is the abbreviation for "Musical Instrument Digital Interface" and is a standardized interface for synthesizers, sequencers, home computers, rhythm devices and so on. Devices equipped with this interface can be operated together. Thus it is possible to control synthesizers with a MIDI interface from a computer such as the ST.

From a hardware standpoint, the MIDI interface consists of a sender and a receiver for asynchronous serial transmission similar to the RS-232 interface. The MIDI transmission takes place at 31.25Kbaud with one start bit, eight data bits and one stop bits. This device therefore requires 320 microseconds to transfer these 10 bits, which corresponds to 3125 bytes per second. This high data rate is required so that real-time applications are possible.

Communication over the MIDI interface takes place with multi-byte messages. The first byte is a status byte. The status byte is denoted by a set seventh bit. The lower 4 bits select one of 16 channels to which the message applies. It is also possible to send messages to all connected devices.

The ST has connections for MIDI IN and MIDI OUT for receiving and sending MIDI information. The optional MIDI THRU connection can be performed through software.

Operating system procedures are available to service the MIDI interface. These functions can output data to and read data from the MIDI interface.

The actual data input and output is interrupt controlled and is carried out by a 6850 ACIA circuit (Asynchronous Communications Interface Adapter). Another operating system function returns the status of the MIDI interface. This allows us to determine if data is present at the MIDI interface or if it is ready for data. An operating system call can also be used to set the size and address of a buffer for temporary storage of the MIDI data received.

6.3 The ST Keyboard

A serious computer needs a solid and reliable keyboard. Factors that determine the suitability of a keyboard include the number of keys, the layout of the keys and the ergonomic organization of the keyboard. The keyboard must not be too high and must be at the proper angle. Only when these things are correct is it possible to work at the computer for a long time without fatigue.


The keyboard on the ST is divided into four sections. The first and largest section is a typewriter keyboard. Two large shift keys switch between lower and upper case or special characters. The control key in combination with other keys creates the non-printable control characters from 0 to 31. There is also an ESC (escape) key and an ALTeRNate key. This makes it possible to create any desired ASCII code. After the ALT key is depressed, three digits representing the ASCII code may be entered. The upper case letters can also be obtained without the shift keys by using the CAPS LOCK key.

The second section of the keyboard is the cursor pad. This consists of four cursor-control keys arranged in a T- shape. The HOME/CLEAR key places the cursor in the upper left-hand corner of the screen or, if used with the shift key, also erases the screen. A character can be inserted at the current cursor position with the "INSERT" key. The keys "HELP" and "UNDO" are also part of this section.

The numeric keypad is located on the right side of the computer. It allows fast entry of numerical data. In addition to the numerals and decimal point, the numeric keypad also has the keys "+", "-", "*", and "/" for the standard arithmetic functions, "(", ")" for mathematical operations and an "ENTER" key for terminating the input.

The ST has ten function keys designated F1 through F10 located above the typewriter section. These keys can be assigned special purposes through software.

Figure 6-6: Keyboard layout



6.3.1 The ST Keyboard processor

From a hardware standpoint, the keyboard actually appears to be an "intelligent" peripheral to the 68000 processor. The keyboard has a single-chip microcontroller, the HD63P01M1 from Hitachi, which has 29 I/O lines, a 16-bit timer and a serial interface built in. This CMOS processor also contains 128 bytes of RAM on the chip. In the prototype of the Atari ST, this processor has a 4K "piggy back" EPROM which contains the operating system for the keyboard processor. In production versions, the processor with a built-in ROM (HD63P01V1) will replace the EPROM.

Communication from the HD63P01M1 takes place over a serial interface already intergrated into the controller. The 68000 uses a 6850 ACIA (Asynchronous Communications Interface Adapter) as its peripheral component. The data transmission takes place at a rate of 7800 baud. Because the serial interface is designed for both input and output, you can "program" the keyboard by sending data to the controller to perform such functions as CAPS LOCK.

The advantage of an intelligent keyboard lies in the fact that special functions can be achieved easily without having to burden the main processor. This single-chip controller frees the 68000 from the task of decoding the keyboard matrix. It is well suited for this task since its has 29 I/O lines. The 68000 gets the code of the key which is pressed just by requesting it from the controller.

This single-chip controller also handles input from the joysticks and/or mouse. The I/O lines that are not used for polling the keyboard are used for this. Two connectors on the right side of the 520 ST are used to attach the joysticks or the mouse. One port is used for the mouse. The four directions and the fire button on the joystick can be polled while the mouse sends the direction of movement or button press. The single-chip processor sends the keyboard, joystick, and mouse data to the 68000 via the serial interface. Since the data causes an interrupt at the 6850, it is made immediately available to the 68000. The character can then be read from the single-chip controller by the interrupt service routine.

The controller is an eight-bit CMOS processor which can be equipped with a 4 or 8K EPROM. It's instruction set is compatible to the 6800 and 6801. It also has bit-oriented instructions for polling individual I/O lines.

6.4 The ST Video Interface

Since the ST does not have a built-in monitor, it must be connected to a video display terminal for operation. The simplest output is to connect the ST to a color or black-and-white television set. The 1040 ST has a built-in RF modulator to produce the necessary signals for the television receiver. But because of its limited screen resolution, a television set is not the optimal output device. Let's look at the other options for screen display.

The ST uses only the graphics mode for text and graphics output.

When creating a video picture, the screen is divided into individual points which can be separately set or cleared. On the screen, a set point appears white and a cleared point black. This assignment can be changed, however. Different colors can be assigned to these two states, for example.


For the sake of simplicity, let's start with the monochrome mode on the ST. In this mode, the screen resolution is 640 points horizontally by 400 points vertically. This gives a total of 256,000 points which can be set or cleared independent of each other.

Each screen point can have one of two different states. In memory, one bit also has two possible conditions: zero or one. This suggests that each screen point be assigned to a bit in memory. In order to achieve the resolution mentioned before, we need 256,000 bits or 32K bytes.

This is referred to as bit-mapped graphics. This 32K must be placed somewhere in RAM. With 512K of RAM, however, this is only one sixteenth of the total RAM.

The ST can do more than monochrome display. In addition to the high-resolution single-color mode with 640x400 screen points, the ST has another mode in which four colors can be displayed simultaneously on the screen. The resolution in this mode is now 640x200 points. How can we display four different colors with only half as many points? To answer this, we must consider how the colors are stored in the 32K of graphics RAM.

Figure 6-7: 640 x 400 Bit-mapped graphics



When one bit is used for a screen point, this bit can take on one of only two conditions, just as the corresponding screen point can be either set or unset. In the color mode, two bits are used for representing a single screen point. These two bits can contain four different values.

0 0	=> 0
0 1	=> 1
1 0	=> 2
1 1	=> 3

A color can be assigned to each of these four values. The video controller is responsible for the assignment and creation of the screen signals as is often referred to as a CRTC (Cathode-Ray Tube Controller).

In addition to the four-color mode which provides 128,000 points of resolution, there is a multi-color mode in which it is possible to display 16 different colors at the same time. By simply extending the method used for the four-color mode, we can use four bits per point to represent 16 colors:

0 0 0 0	=> 0
0 0 0 1	=> 1
0 0 1 0	=> 2
0 0 1 1	=> 3
0 1 0 0	=> 4
0 1 0 1	=> 5
0 1 1 0	=> 6
0 1 1 1	=> 7
1 0 0 0	=> 8
1 0 0 1	=> 9
1 0 1 0	=> 10
1 0 1 1	=> 11
1 1 0 0	=> 12
1 1 0 1	=> 13
1 1 1 0	=> 14
1 1 1 1	=> 15

Since we have 256,000 bits of memory and we need four bits for each point, we can control 64,000 points. This gives us a resolution of 320 points horizontally and 200 vertically.

So the ST can display either four or sixteen colors at a time. These colors may be any one of 512 colors. These colors are combinations of the primary colors red, green and blue. Each of these three primary colors can be used

in one of eight brightness levels. This results in $8 * 8 * 8 = 512$ combinations. With these colors available, practically all color gradations can be represented. White results from a superimposition of all three primary colors. If the primary colors are used only in brightness levels, corresponding gray levels are produced. The combination of red and green gives yellow, red and blue make purple (magenta), and blue and green make cyan. All other colors can also be formed with a partial mix of two primary colors.

To achieve the best picture quality, you should use the RGB output. RGB is an abbreviation for Red Green Blue. With this interface, the signals for the individual colors are available separately and can be used directly for controlling the cathode rays in the monitor. The ST has such an RGB interface for direct connection of an RGB monitor.

A great loss in quality happens if you connect the ST to a television set. Here the video signal is first modulated into a high-frequency signal and then demodulated in the receiver. The signal must be further divided into the three primary colors. In addition, televisions are not designed to display pictures with as high a resolution as the ST generates.

The ST can also be connected to a high-resolution monochrome monitor to make full use of the high-resolution in the single-color mode.

With its video interfaces, the ST offers something for differing tastes in quality and price range. Atari will offer a high-resolution 12" monochrome and a 12" RGB color monitor for the ST.

Figures 6-8 and 6-9 show screen photos of the Atari ST which demonstrates the high resolution of 640x400 points.

Figure 6-8: HI-Resolution ST Screen Photo



Figure 6-9: HI-Resolution ST Screen Photo



Chapter 7

The Operating System (TOS)

- 7.1 The BIOS and BDOS**
- 7.2 BDOS Functions**
- 7.3 The BIOS**
- 7.3.1 The XBIOS**

7 The Operating System (TOS)

Now that we've become a bit more acquainted with the hardware of the ST, let's dig deeper into the operating system. First let's again clarify what an operating system is, what tasks it has to perform, and what demands are placed on a modern operating system.

The capable hardware of the ST is of no use without a program to direct the processor and peripheral components. The operating system is the "director". The operating system makes it possible for the computer to recognize a press on a key; to construct a command from multiple key presses; to carry out the commands that are entered; to send data to and receive data from peripheral devices that result from these commands and other functions as well.

For any computer, the primary input and output devices are the keyboard and the display screen. They make it possible to work with the computer interactively - the computer responds immediately to keyboard input with results on the screen. If the results are to be documented, the output can be routed to a printer. The operating system controls the printer so that it produces the desired results.

The ST can support two joysticks and a mouse as additional input devices. The mouse is an input device which can simplify the input of data between the user and computer. To make such devices easy to use, the operating system must provide procedures to service these devices.

In order to exchange data with other computers, you can use a telephone, a modem, and a computer with a serial interface. Using this configuration you can connect to almost any computer in the world. These may be simple home computers or complex mainframes which manage huge data banks. Controlling the ST's serial interface also falls to the operating system.

The disk drive is a necessity for most computers. Recently hard disk drives, such as the 3 1/2" drives from Atari, have become available at a price which makes them a reasonable alternative to floppy drives. In most instances, the operating system is contained in the computer's memory while user programs such as word processing, data management, or high-level programming languages are kept on the mass storage and loaded into memory by the operating system when needed. After the program is loaded, the operating system passes control to the user program.

An important task of the operating system is that of moving programs and data between the computer and the mass storage. Often the operating system itself is stored on the mass storage device and is loaded into the computer when it is turned on. The name "Disk Operating System", or DOS for short, is used to name an operating system whose main task is the management of the disk drive.

But an operating system does more than simply load programs from the mass storage devices and then start them. It must also perform the input and output operations for user programs.

In the hardware description of the ST we became acquainted with the peripheral components responsible for data input and output. If a user wishes to output text to a printer from within a program, the computer must know exactly which peripheral component the printer is connected to. Even if this is known, it is not enough to simply send the data to the peripheral component; the data that is output to the printer must also follow a certain *protocol*. In the description of the printer interface we saw that this is done with the STROBE and BUSY lines.

In general, a protocol is a set of rules which determines how data is sent between two devices. There is a protocol for serial interfaces too, such as the XON/XOFF protocol described earlier. Since input and output is performed by every program, it doesn't make sense for the user to program these things over and over again for each program. Furthermore, to program at this level requires that the programmer have a very detailed knowledge of both hardware and machine language.

So the operating system performs most of these tasks. The operating system contains ready-to-use routines called input/output routines. By using these routines you can output text to the screen, read characters from the keyboard, send or receive data over the serial interface, produce hardcopy on the printer or read data from or write data to a diskette.

To use these functions, the programmer simply calls the appropriate operating system routine. Since the operating system usually has to input or output data, the programmer must inform the operating system the location of the data and the type of transfer that is desired.

How should an operating system routine be designed? Before answering this question, let's pause a few minutes to talk about the early days of small computers.

When the first personal computer appeared on the market ten years ago, its price was determined largely by the cost of the hardware. The cost of memory for example, was relatively expensive. This was the reason that most of the early computer had small amounts of memory - either 4K or 8K of RAM. Many users couldn't afford to buy more. But advances in production techniques and volume sales have reduced the cost per kilobyte to less than 25 cents. Within a year, 1M (million) bit chips will most likely be available and most personal computers will have at least 1M-bytes of RAM. By replacing 256K-bit chips with 1M-bit chips the ST can reach its 4M-byte internal memory limit.

Similarly, advances have been made in the manufacturing of mass storage devices that have reduced the prices of both floppy and hard disk drives.

In the past, the cost of the hardware components accounted for 80 to 90% of the price of a personal computer. As a result of these tremendous drop in component cost, the price of the more recent computers have been steadily sinking.

There are some differences between the software and hardware costs. Developing software has been and still is labor intensive. As such, the many man-months or man-years that it takes to develop complex software is expensive. In order to maximize the software development investment, most developers try to design their software so that it can run on as many computers as possible.

One operating system which was developed with the idea of being used on many different computers, is called CP/M (Control Program for Microcomputers). Developed by a company named Digital Research to work on the Intel 8080 processor, it was later adapted to work with the Zilog Z-80 processor. It is written in such a way as to be able to work on computers developed by different manufacturers. The ideas learned from the development of CP/M were used in the development of TOS.

At the start of 1985 Digital Research began developing an operating system designed for 68000 computers called GEMDOS. GEMDOS includes the hardware-independent functions of TOS. These functions allow the programmer to easily control the input/output functions of the computer.

To ensure that program developers would be able to master this new operating system quickly the functions in GEMDOS are very similar to the functions of MS-DOS. The function numbers used correspond to those of MS-DOS. MS-DOS is the operating system used in the IBMs and compatibles. A large group of programmers are already familiar with this operating system. Not all of the MS-DOS functions are included in GEMDOS, in the area of file management only the UNIX compatible functions are implemented. UNIX is an operating system that runs on much larger computers and some people believe that this operating system will eventually become the standard.

7.1 The BIOS and BDOS

To make it easy to adapt one operating system to the differing computer hardware, they were divided into three software components, each functionally separate from one another. These components are:

BIOS	BASIC INPUT OUTPUT SYSTEM
BDOS	BASIC DISK OPERATING SYSTEM
CCP	CONSOLE COMMAND PROCESSOR

The BIOS, or basic input output system, performs the actual interfacing to the hardware. This part of the operating system includes the routines for screen output, reading the keyboard, printer output and reading and writing individual sectors on floppy or hard disk.

The BIOS is hardware specific and must be rewritten to implement the operating system on a new computer. The BIOS has a preset number of simple input and output operations to perform. At the beginning of the BIOS program is a *jump table* for all of its operations. This table contains the memory location of the routine which performs a given operation. The routines themselves are adapted specifically for each different computer.

The BDOS, or basic disk operating system, is responsible for the logical organization of mass storage data. It manages data at the file level.

One of its main jobs is to manage the disk storage space. When you create a new file, for example, the BDOS searches for unused space on the disk. It selects the track and sector at which the data is stored. The user only needs to know the name of the file. To help manage the disk space, the BDOS use a directory containing the file names, type, size and location of the file. So when you later want to access the data in the file, the BDOS knows where to find it.

To reiterate, the BIOS is in charge of the physical management of the peripheral devices while the BDOS is responsible for the logical management of the disk drive. With help from the BDOS, programs and data can be stored on the disk and read back again.

The third component of the operating system, is the CCP or console command processor. It is the part of the operating system that is responsible for handling the main line of communication between the user and

operating system from the keyboard. When CP/M starts up, the CCP responds with the following output on the screen:

A>

The A means that the first disk drive is selected. Under CP/M the drives are designated with letters which range from A to P. Therefore it is possible to connect up to 15 drives. The character > indicates that the operating system is now waiting for the user to input a command. We can now send a command to the CCP. To do this, we type the command on the keyboard. The characters are echoed on the screen as the user types them. When a command has been entered and terminated by pressing the RETURN-key, the CCP analyzes it.

As you can see this is a very cryptic operating system in which the user has to remember a large command set to use the computer. The GEM interface is a major improvement to the CCP. Compare the ST screen with the >A: prompt and decide which operating system a first time user would be more comfortable with.

The actual internal operation of TOS is very similarly to CP/M 2.2 for the eight-bit computers. It is also divided into the three components: the BIOS, BDOS and CCP. But there are a few major differences:

- * The maximum memory size for CP/M 2.2 is 64K but for TOS the limit is 4MB.
- * For CP/M 2.2, the operating system is loaded into memory from diskette when the computer is turned on. On the ST, however, the TOS is stored in ROM (early models loaded TOS from disk.) TOS in ROM has several advantages. It is available immediately after the computer is turned on since it does not have to be loaded from disk first. Second and more importantly, the user loses no memory space to the operating system.
- * For CP/M 2.2, the operating system is always placed at the end of memory. The TOS operating system may be placed anywhere in memory.

7.2 BDOS Functions

A user program can communicate with the peripheral devices via BDOS calls. Logical input and output devices are defined in the BDOS which are first assigned to physical devices in the BIOS. The following input devices are supported by the BDOS:

CON: STANDARD INPUT, USUALLY KEYBOARD
AXI: AUXILIARY INPUT, USUALLY SERIAL INTERFACE

Three output devices are available:

CON: STANDARD OUTPUT, USUALLY SCREEN
LST: LIST DEVICE, USUALLY PRINTER
AXO: AUXILIARY OUTPUT, USUALLY SERIAL INTERFACE

If a service is required of the BDOS by a user program, the BDOS is accessed with a TRAP instruction. The TRAP instruction causes the 68000 to execute an exception handling routine. When the 68000 encounters this instruction, it branches to a routine for handling this condition. Since the BDOS has many functions available, the number of the BDOS function is passed in register D0 of the 68000. Based on this number, a branch is made to the routine which performs the desired function. The additional registers of the 68000 are used to pass parameters for the functions. If a character is to be output to the screen, for instance, it is passed in register D1. On the other hand, if a character is to be read from the keyboard, a call to this BDOS returns the character in register D0.

Here's an example. To print the character X on the screen, a call to the appropriate BDOS function (number 2) might look like this:

```
002000 123C0058      MOVE.B #'X',D1 ; CHARACTER
002004 303C0002      MOVE.W #2,D0  ; BDOS #
002008 4E42          TRAP #2    ; CALL
```

An area called the I/O-byte is defined in the BIOS. Its purpose is to assign logical devices to the physical devices. The I/O-byte is divided into four sets of two bits each. Each of these four groups is associated with a logical device.

I/O-Byte								
dev a		dev b		dev c		dev d		
bit	0	1	2	3	4	5	6	7

Four physical devices can be assigned with these two bits. The bit pattern "00" selects the standard device while a different bit pattern selects an alternate device. This makes it possible to redirect output from the screen to the printer, for example. In the same manner it's possible to reroute input from the serial interface, perhaps through a modem, instead of the keyboard.

In order to give the novice and experienced programmers an overview of what functions can be executed by the BDOS, we will list them individually. In each case the function number is placed in register D0 before calling BDOS.

NUMBER FUNCTION

0 **TERM**

This function is used to return control back to the calling program. For applications started from the desktop, program control is returned to the desktop.

1 **CONSOLE INPUT**

This function inputs a character from the console. Control characters such as CR (carriage return) or LF (line feed) are also returned. This function waits until a character is entered at the console.

2 **CONSOLE OUTPUT**

This function outputs a character to the console. The ASCII code for the character to be output is placed in register D1.

3 **SERIAL INPUT**

This function is like function 1 except that the character is input from the serial channel. This function also waits until a character is received.

4 **SERIAL OUTPUT**

This function is similar to function 2 except that the character is output to the serial port. The ASCII code for the character to be output is placed in register D1.

5 LIST OUTPUT

This function outputs a character to the logical device LST:. The ASCII code for the character is placed in register D1. This device is normally assigned to the Centronics parallel interface, to which a printer is connected.

6 DIRECT CONSOLE INPUT/OUTPUT

This function is intended for special purposes where the normal BDOS functions do not work. Using this function, output to the screen cannot be stopped with CTRL-S and cannot be redirected to the printer with CTRL-P.

7 CONSOLE INPUT WITHOUT ECHO

This function is similar to 1 except that the character received is not displayed on the screen.

8 CONSOLE INPUT WITHOUT ECHO

This function is the same as 7. This was done for MS-DOS compatibility.

9 OUTPUT STRING

This function is used to output a string of characters. The string may be stored anywhere in memory. Register D1 contains the address of this string. The end of the string is terminated with a zero byte.

10 READ CONSOLE BUFFER

This function reads an entire input line from the keyboard. The BDOS waits until the input is terminated with a RETURN.

11 GET CONSOLE STATUS

This function informs the caller if a character has been input from the console.

14 SELECT DISK

This function selects a drive. The drive number is passed to the BDOS in register D1.

16 CONSOLE OUT STATUS

This function returns the status of the console in register D.

17 PRINTER OUTPUT STATUS

This function returns the status, the condition of the centronics interface.

18 AUXILIARY INPUT STATUS

This function determines if a character is available from the receiver of the serial interface.

19 AUXILIARY OUTPUT STATUS

This function gives information about the state of the serial bus.

25 RETURN CURRENT DISK

This function returns the current default drive.

26 SET DMA ADDRESS

This function sets the address of the buffer in which the data from the disk will be written.

32 SUPER

This function selects supervisor or user mode.

42 GET DATE

This function returns a coded value in D0 representing the date.

43 SET DATE

This function allows the date to be set.

44 GET TIME

Returns the current time from the GEMDOS clock.

45 SET TIME

Sets the GEMDOS clock.

47 GET DISK TRANSFER ADDRESS

This function returns the address of the current disk transfer buffer.

48 GET VERSION NUMBER

This function returns the current version number of GEMDOS.

- 49 KEEP PROCESS**
Similar to the function **TERM** except the memory is not cleared when the program is ended.
- 56 GET DISK FREE SPACE**
This function enables you to calculate the number of free bytes available on a diskette.
- 57 MAKE DIRECTORY**
This function creates a new folder (directory).
- 58 REMOVE DIRECTORY**
This function allows for the deletion of empty folders.
- 59 CHANGE DIRECTORY**
This allows access to different folders (sub-directories).
- 60 CREATE**
This function creates a file on the disk drive.
- 61 OPEN**
To process existing files they are opened with this function.
- 62 CLOSE**
Files are closed with this function.
- 63 READ**
Opening and closing files is only part of the process. The data in the file must be read into the computer with this function.
- 64 WRITE**
Data is written to a file with this function.
- 65 UNLINK**
Files that are no longer needed can be removed with this function.
- 66 LSEEK**
This UNIX-compatible command is used for access to relative files.
- 67 CHANGE MODE**
This function is used to change the parameters of a disk file.

- 69 DUP**
Different devices can be assigned a file handle number with this function.
- 70 FORCE**
Input or output can be redirected with a call to this function.
- 71 GETDIR**
This function returns the current pathname.
- 72 MALLOC**
This is a memory allocation function used to reserve memory.
- 73 MFREE**
Memory allocated with MALLOC is released with this function.
- 74 SETBLOCK**
This function is used to reserve the actual memory requirements of a program and release the remaining memory.
- 75 EXEC**
This function permits loading and chaining of programs.
- 76 TERM**
This terminates a program with a programmer-defined return value returned to the calling program.
- 78 SFIRST**
This function is used to see if a file with the given name is present in the directory.
- 79 SNEXT**
This function can be used to see if there are any other files on the disk which match
- 86 RENAME**
This is used to rename files on the disk.
- 87 GSDTOF**
The date and time of creation of a file can be read or set with this function.

7.3 The BIOS

The software interface between GEMDOS and the hardware is the BIOS (Basic Input Output System). The BIOS for the ST has additional functions which support the special hardware attributes of this machine. The standard BIOS functions are called with TRAP #1. The extended (XBIOS) functions are called with TRAP #14.

Normally user programs do not use BIOS calls. Instead they use the logical functions provided by the BDOS.

The following list contains a brief description of the BIOS calls.

NUMBER	FUNCTION
--------	----------

- | | |
|---|--|
| 0 | GETMPB
This function returns the address of the memory parameter block. |
| 1 | BCONSTAT
This function has the same task as the corresponding BDOS call. A value of ff (Hex) is returned in register D0 if a character is ready, otherwise a zero is returned. |
| 2 | CONSOLE INPUT
A character input from the console device is returned in register D0. |
| 3 | CONSOLE OUTPUT
The ASCII code contained in register D1 is sent to the console. |
| 4 | READ-WRITE A DISK SECTOR
This function serves to read and write sectors on the disk. |
| 5 | SET EXCEPTION VECTORS
This function allows one of the exception vectors of the 68000 processor to be set. |

- 6 TIC CALCULATE**
This function returns the number of milliseconds between two system timer calls.
- 7 GET BIOS PARAMETER BLOCK**
A pointer to the BIOS Parameter Block is returned.
- 8 OUTPUT DEVICE STATUS**
This function checks if the output device is ready to output the next character.
- 9 INQUIRE MEDIA CHANGE**
Determines if the disk was changed since its last access.
- 10 INQUIRE DRIVE STATUS**
This function returns a bit vector which contains the connected drives.
- 11 INQUIRE? CHANGE KEYBOARD STATUS**
With this function it is possible to determine or change the status of the special keys on the keyboard.

7.3.1 The XBIOS

On the next pages you find the eXtended BIOS (XBIOS) functions. They are used for handling the special hardware features of the Atari ST hardware.

NUMBER	XBIOS	FUNCTIONS
--------	-------	-----------

- | | | |
|---|---------------------------------------|---|
| 0 | Initialize mouse | This function initializes the routines for mouse processing. |
| 1 | Save memory space | Memory can be reserved before the operating system is initialized with this function. |
| 2 | Return screen RAM base address | This function returns the base of the physical screen RAM. |
| 3 | Set logical screen base | This function sets the address used for all output functions. |
| 4 | Return screen resolution | The screen resolution currently active can be determined with this function. |
| 5 | Set screen parameters | This function sets various screen parameters. |
| 6 | Set color palette | New color palettes may be loaded with this function. |
| 7 | Set color | This function allows one color to be changed. |
| 8 | Read diskette sector | This function reads one or more sectors from the diskette. |
| 9 | Write diskette sector | One or more disk sectors can be written to the disk using this function. |

- 10 Format diskette**
This function will format a track on the diskette.
- 11 Unused**
- 12 Write string to MIDI interface**
Strings may be output to the MIDI interface using this function.
- 13 Initialize MFP format**
This function initializes an interrupt routine in the MFP.
- 14 Return record buffer**
This returns a pointer to a buffer data record for an input device.
- 15 Set RS-232 configuration**
You can easily configure the RS-232 interface parameters with this function.
- 16 Set keyboard table**
This function allows the creation of a new keyboard table.
- 17 Return random number**
A 24 bit random number is returned with this function.
- 18 Produce boot sector**
You can create a boot sector with this function, possibly to load another operating system.
- 19 Verify diskette sector**
This function verifies one or more sectors on the disk.
- 20 Output screen dump**
A hardcopy of the screen is output to a printer.
- 21 Set cursor configuration**
You can set the cursor to steady, flashing or invisible with this function.
- 22 Set clock time and date**
The clock time and date are set with this function.

- 23 Return clock time and date**
This function returns the current clock time and date.
- 24 Restore keyboard table**
This function restores the standard keyboard layout.
- 25 Transmit command to keyboard**
Commands can be sent to the intelligent keyboard processor with this function.
- 26 Disable interrupts on MFP**
Interrupts on the MFP may be selectively disabled with this function.
- 27 Enable interrupts on MFP**
This function is used to re-enable interrupts on the MFP.
- 28 Access GI sound chip**
Access to the registers of the General Instruments sound is possible with this function.
- 29 Reset port A of GI sound chip**
A bit of port A of the sound chip can be set with this function.
- 30 Clear port A of GI sound chip**
A bit of port A of the sound chip can be cleared with this function.
- 31 Start MFP timer**
You can start a timer in the MFP with this function.
- 32 Set sound parameters**
This function allows for easy sound processing.
- 33 Set printer configuration**
This function is used to read or change the printer configuration.
- 34 Return keyboard vector table**
This function returns a pointer to a table that contains the addresses of the routines which process the data from the keyboard processor.

- 35 Set keyboard repeat rate**
The keyboard repeat rate can be set with this function.
- 36 Output block to printer**
This function is used by function 20 (scremp).
- 37 Wait for video**
This function waits for the next picture return, it can be used to synchronize graphic outputs.
- 38 Set supervisor execution**
This function allows routines to be run in supervisor mode.
- 39 Disable AES**
The AES can be disabled with this function, provided it is not in ROM.

Chapter 8

Inside GEM

- 8.1 Virtual Device Interface**
- 8.2 Application Environment Services**

8. Inside GEM

GEM is an acronym for Graphics Environment Manager. The complete GEM software package is a graphics-oriented user interface designed to simplify the operation of the computer.

There are two major components to GEM. They are the VDI (Virtual Device Interface) and the AES (Application Environment Services).

The ST was designed with simplicity in mind. To make working with the ST as easy as possible, all user interaction with the operating system and application are meant to be handled in an identical fashion. This considerably shortens the training time to learn a new application, since the basic operations of using the mouse, icons, pull-down menus and windows remains the same. In addition, software designed for GEM works independent of the computer hardware configuration. So if the hardcopy device for the ST is changed from a dot matrix printer to a more expensive plotter, the output is easily produced.

This compatibility is achieved through the *virtual device interface*. The idea behind this is a software *driver* which is responsible for the output of data to hardware peripheral.

All input and output for the ST are sent to the VDI and not directly to the operating system. The VDI accepts incoming data and reproduces it on the screen, on the printer, on the plotter, etc.

Another component of GEM is AES or application environment services. The AES monitors the input devices such as reading the keyboard, evaluating the movement of the mouse, and checking the mouse button. The AES also controls output formatting such as windows, pop-down menus and icons.

These two parts of GEM, the VDI and the AES, along with the TOS form the block of system-specific software. Included in this system is the DESKTOP which can call the various utility programs, the TOOLS. These might include an icon editor with which you could create your own icons.

8.1 Virtual Device Interface

The VDI itself, consists of two parts: namely the GDOS, an operating system for controlling the graphics-oriented output devices, and the driver which contains the information concerning the character sets (or fonts) to be used.

This differentiation suggests one essential difference between GEM and earlier operating systems. On conventional computers, there are two types of screen displays: a *TEXT* mode and a *GRAPHICS* mode. In the *TEXT* mode display, one byte of the screen memory is interpreted as a character according to the assignments of the ASCII codes. In the *GRAPHICS* mode, on the other hand, when between 50,000 and 200,000 points are addressed, a single bit represents only one of these points.

As a result, the size of the screen memory is dependent on the screen display mode. Further, trying to combine text and graphics on the same display are made more difficult.

Using GEM, there are no provisions for a pure *TEXT* mode; the information is always displayed in the *GRAPHIC* mode.

To display an alphanumeric character at a specific location, the bits in the corresponding screen memory are set to the pattern of the desired character. This is called *bit-mapped* character representation. Since any pattern can be displayed, the user is free to use any set of alphanumeric patterns. Such a set of patterns is called a *character set*. Different character sets may be stored on diskette and used in place of a standard character set. A common use for user-defined character sets is to display foreign characters or to substitute a new *font*. Thus alphanumeric characters are handled just like graphic patterns in a bit-mapped system.

Here's another example of the drawback of a separate *TEXT* and *GRAPHIC* display mode. Suppose you want to display a square box and a perfect circle on the screen and then get a hardcopy of them on the computer's dot matrix printer. Using most computers, the square is reproduced as a rectangle and the circle as an ellipse. This is due to the differences in point representation on the screen and the printer.

The VDI can overcome this problem. The VDI recognizes two different coordinate systems, namely:

- * NDC - Normalized Device Coordinates, and
- * RC - Raster Coordinates.

Raster coordinates correspond exactly to the physical capabilities of a display device. For the ST display screen, this is either 320 X 200 or 640 X 400 points.

Normalized device coordinates correspond to an idealized screen surface. Using normalized coordinates, a quadrilateral that appears square on the screen is made to look square when printed as hardcopy. For a coordinate system of say, 600 x 200, the height to width relationship between the coordinates on the screen is 1:1.8. That is, one screen pixel in the horizontal direction has the same dimensions as 1.8 pixels in the vertical direction.

With the VDI, a software *device driver* carries out the necessary conversions for NDC. The output appears in the same proportions on all peripheral devices.

To do this, GEM use *meta-files*. These files contain information required to exchange data among peripheral devices and individual GEM applications. In a meta-file, the physical size of the output device and the coordinate system for the data to be transmitted is defined.

8.2 Application Environment Services

The AES is composed of several components. The most important are the dispatcher, the screen manager and the subroutine library.

The *dispatcher* is responsible for handling multiple applications that are able to run simultaneously. Since the 68000 processor is so fast compared to the eight-bit processors, it is possible for the ST to perform more than one task at a time. This is called *multitasking*. Later versions of GEM may support multitasking. Currently GEM is limited to running one application, three desktop accessories or background processes (printer spooler) and the AES.

The subroutine library contains the appropriate routines for manipulating the windows, handling the mouse, and so on. The screen manager assumes control of the mouse once it is outside of a window.

The use of the AES function calls will greatly speed the development of GEM software for the ST. Because the AES is extremely powerful and therefore a complicated subject we will not present anymore details of the powerful system. Those interested should get Abacus Software's Atari ST GEM Programmer's Reference for complete details on the AES and VDI function calls.

Chapter 9

CONCLUSION

9. CONCLUSION

The Atari ST is truly a computer that features "**POWER without the PRICE**". The TOS operating system is a great leap in the user interface of computer systems. The first time user should be able to use the ST with almost no training. The free software included with the system would cost hundreds of dollars on other systems.

The decision by Atari to use industry standard interfaces insures that a wide variety of peripheral devices will be available of the ST today and tomorrow. The price of this computer will certainly make it one of the most popular home and business computers of this century.

We hope that this brief introduction has given the novice and experienced user enough information to intelligently decide if this is the computer system for their application.

Please remember that this is a presenting book and by the time it is published the excellent development team of Atari will probably have introduced many new software and peripheral devices. Watch for many new and exciting applications from the new ATARI for this amazing computer.

INDEX

A

Addressing modes, 98
ADSR, 133
Amplitude control, 133
ASCII, 120, 124, 135
Asynchronous, 120, 123
Asynchronous Communications Interface Adapter, 134, 137
Attack-Decay- Sustain-Release, 133
AY-3-8910, 132

B

B/W Monitor, 112
Bank-switching, 86
BASIC, , 79, 85
Baud rate, 123-126
BDOS, 151, 153
BIOS, 151, 159
Bit-mapped graphics, 139
Bus error, 105
BUSY, 121

C

CAPS LOCK, 135
Cathode-Ray Tube Controller, 140
CCP, 151, 152
Centronics parallel interface, 112, 120
Checksum, 115
CHK, 105
Clicking, 15
CMOS, 137
COBOL, 85
Color monitor, 112
Colors, 138
Composite video interface, 141
CRC (Cyclic Redundance Check), 115
Cursor-control keys, 135

D

D/A converter, 132
Data integrity, 118
Data register, 93
Data transmission rate, 117, 118, 137
Dialog box, 19
Digital Research, 150
Disk controller, 114
Disk drives, 114-117
DMA, 90, 114, 116
DOS, 148

E

ENTER, 137
EPROM, 137
Expansion, 111

F

FORTTRAN, 85
Four-color mode, 140
Function keys, 135

G

GEM, (Graphics Environment Manager) , 167

H

HARD DISK, 116, 117
HD63P01M1 from Hitachi, 157
HD63P01V1, 137
Head crash, 118
HELP, 135
High-level languages, 61, 85

I

I/O interface, 118
IBM PC, 87
IBM-370, 90
INSERT, 135
Interface, 116
Interrupt control, 88
Interrupt controller, 126, 128
Interrupt priority, 104
Interrupt sources, 130

J

JOYSTICK, 112, 137, 147
Jump table, 147

K

KEYBOARD, 112, 137

L

LEA (Load Effective Address), 103
LINK, 105
Linker, 106
Long-word, 93

M

Mainframe, 85
MFP 68901, 126, 130
MIDI, 112, 134
Mnemonic, 95
Monochrome mode, 120, 138
Monochrome monitor, 141
Motorola, 88
MOUSE, 10, 112, 147
Multi-color mode, 138
Multi-user, 104

N

NMI, 128
Noise generator, 130, 133
Non-maskable interrupt, 104, 128

O

Operating system, 104, 147

P

Parallel interface, 120
Program counter, 92
Programmable mixer, 132
Protocol, 148
PSG, 132

R

RAM, 109
Recursive calls, 106
RESET command, 104
Resolution, 138
RGB, 112, 141
ROM, 111
ROM cartridge, 109
RS-232, 122, 123
RTS, 123

S

Serial interface, 122, 123
Sound generator, 132
SSP (system stack pointer), 104
Stack, 104
Status register, 92
Supervisor mode 103, 104
Synchronous, 122, 123
Synthesizer, 112, 134

T

Television, 111
Timers, 126
Tone generator, 132, 133
TOS (Tramiel Operating System), 4
Tpi (tracks per inch), 116
Track density, 116
Two's complement arithmetic, 97

U

UNDO, 136
UNLK, 106
USART, 130
User mode, 103, 104
USP (user stack pointer), 104

V

VIDEO, 112
Volume levels, 132

W

Wave forms, 132, 133
WD 1770 from Western Digital, 114

X

XOFF, 123
XON, 124
XON/XOFF, 148

Y


YM 2149 from Yamaha, 132

Z


Z-80, 85, 87, 150
Zilog, 150

ATARI[®] ST[™]

REQUIRED READING



A fantastic collection of useful programs and information for the ST. Complete programs include: super-fast RAM disk; time-saving printer spooler; color print hardcopy; plotter output hardcopy. Explains BASIC commands to access GEM using VDISYS and GEMSYS and describes resource files with examples. Manipulate text output (size, rotation, bold, etc.) Change line types (thickness, endpoints, etc.) Mixing machine language with BASIC or C programs. Save software dollars with these tricks and tips. 250 pages \$19.95



For the serious programmer in need of detailed information on the GEM operating system. Written especially for the Atari ST with an easy-to-understand format that even beginners will be able to follow. All GEM routines and examples are written in C and 68000 assembly language. Covers working with the mouse, icons, Virtual Device Interface (VDI), Application Environment Services (AES) and the Graphics Device Operating System. Required reading for the serious programmer interested in understanding the ST. 450 pages. \$19.95



MACHINE LANGUAGE
Program in the fastest language for your Atari ST. Learn the 68000 assembly language, its numbering system, use of registers, the structure & important details of the instruction set, and use of the internal system routines. 280pp \$19.95

ST INTERNALS
Essential guide to learning the inside information on the ST. Detailed descriptions of the sound & graphic chips, internal hardware, the I/O ports, system addresses, more. Fully documented BIOS assembly listing. An indispensable guide. \$19.95

GRAPHICS & SOUND
A comprehensive handbook showing you how to create fascinating graphics and surprising music and sound from the Atari ST. See and hear what sights and sounds that you're capable of producing from your Atari ST. \$19.95

LOGO
Take control of your Atari ST by learning LOGO—the easy-to-use, yet powerful language. Topics covered include structured programming, graphic movement, file handling and more. An excellent book for kids as well as adults. \$19.95

PEEK & POKES
Enhance your programs with the examples found within this book. Explores using the different languages BASIC, C, LOGO and machine language, using various interfaces, memory usage, reading and saving from and to disk, more. \$16.95

PRESENTING THE ST
Gives you an in-depth look at this sensational new computer. Discusses the architecture of the ST, working with GEM, the mouse, operating system, all the various interfaces, the 68000 chip and its instructions, LOGO. \$16.95

The Atari logo and Atari ST are trademarks of Atari Corp.

Abacus Software

P.O. Box 7219 Grand Rapids, MI 49510 - Telex 709-101 - Phone (616) 241-5510

Optional diskettes are available for all book titles at \$14.95

Call now for the name of your nearest dealer. Or order directly from ABACUS with your MasterCard, VISA, or Amex card. Add \$4.00 per order for postage and handling. Foreign add \$10.00 per book. Other software and books coming soon. Call or write for free catalog. Dealer inquiries welcome—over 1400 dealers nationwide.

PRESENTING THE



First, there was the fabulously successful VIC-20. Then came the record-breaking Commodore-64.

Now Jack Tramiel has launched his third home computer, the ATARI ST.

The ST promises to shatter all existing price-performance barriers and become a leader in the home-computer market.

This book, *Presenting the Atari ST* gives you an in-depth look at this sensational new computer that promises to bring you...**"Power without the Price."**

Presenting the Atari ST not only provides the reader with a summary of the capabilities and features of this fascinating new machine, but serves as a good source of information for the prospective buyer.

ISBN 0-916439-33-X

The ATARI logo and ATARI ST are trademarks of Atari Corp.

A Data Becker book published by

You Can Count On
Abacus  **Software**

P.O. Box 7219 Grand Rapids, MI 49510 • Telex 709-101 • Phone 616/241-5510